arXiv:1703.01421v1 [stat.ME] 4 Mar 2017

l_0-ESTIMATION OF PIECEWISE-CONSTANT SIGNALS ON GRAPHS

ZHOU FAN AND LEYING GUAN

ABSTRACT. We study recovery of piecewise-constant signals over arbitrary graphs by the estimator minimizing an l_0-edge-penalized objective. Although exact minimization of this objective may be computationally intractable, we show that the same statistical risk guarantees are achieved by the α-expansion algorithm which approximately minimizes this objective in polynomial time. We establish that for graphs with small average vertex degree, these guarantees are rate-optimal in a minimax sense over classes of edge-sparse signals. For application to spatially inhomogeneous graphs, we propose minimization of an edge-weighted variant of this objective where each edge is weighted by its effective resistance or another measure of its contribution to the graph's connectivity. We establish minimax optimality of the resulting estimators over corresponding edge-weighted sparsity classes. We show theoretically that these risk guarantees are not always achieved by the estimator minimizing the l_1/total-variation relaxation, and empirically that the l_0-based estimates are more accurate in high signal-to-noise settings.

1. INTRODUCTION

Let $G = (V, E)$ be a known (undirected) graph, with vertices $V := \{1, \dots, n\}$ and edge set E. At each vertex $i \in \{1, \dots, n\}$, an unknown signal value $\mu_{0,i}$ is observed with noise:

$$Y_i = \mu_{0,i} + \varepsilon_i.$$

We assume $\varepsilon_1, \dots, \varepsilon_n \overset{iid}{\sim} \text{Normal}(0, \sigma^2)$ and G is fully connected. This paper studies the problem of estimating the true signal vector $\mu_0 := (\mu_{0,1}, \dots, \mu_{0,n})$ based on observing $Y := (Y_1, \dots, Y_n)$, when μ_0 is (or is well-approximated by) a piecewise-constant signal over G. Informally, this will mean that the set of edges $\{i, j\} \in E$ where $\mu_{0,i} \neq \mu_{0,j}$ is a small subset of all edges, E.

Examples of this problem occur in a number of applications:

- **Multiple changepoint detection.** Here, G is a linear chain with n vertices and $n - 1$ edges, which identifies a sequential order to the observations. The signal μ_0 is piecewise constant in the sense $\mu_{0,i} \neq \mu_{0,i+1}$ for a small number of changepoints i. Recovery of μ_0 and similar problems arise in genomics [BBM00, OVLW04], proteomics [AL89], and finance [AIL99, AG02].
- **Image segmentation.** Here, G is a 2-D (or 3-D) lattice graph, and μ_0 corresponds to the pixels (or voxels) of a digital image. The assumption of piecewise-constancy implies that μ_0 has regions of approximately constant pixel value; such images often arise in medical applications [ZBS01, PXP00].
- **Anomaly identification in networks.** Here, G represents a real-world network. The true signal μ_0 indicates locations of anomalous clusters of vertices, for example representing individuals infected by a disease or a computer virus. Piecewise-constancy of μ_0 reflects the assumption that the anomaly spreads along the network connections and that observations for groups of infected individuals exhibit similar behavior.

DEPARTMENT OF STATISTICS, STANFORD UNIVERSITY

E-mail addresses: zhoufan@stanford.edu, lguan@stanford.edu.

ZF was supported by a Hertz Foundation Fellowship and an NDSEG Fellowship (DoD, Air Force Office of Scientific Research, 32 CFR 168a).

Literature concerning this and related problems is extensive: For example, the multiple changepoint problem was studied in [CZ64, Yao84, BH93], the image segmentation problem in [GG84, Bes86], and similar detection/testing problems for general graphs in [ACCHZ08, ABBDL10, ACCD11, SKS13]. The estimation problem we study was also considered recently in [HR16, WSST16, HMPSST16]. We discuss some connections of our work to the existing literature in Section 1.1.

The focus of our paper is the "l_0-denoising" approach to this problem: We consider estimating μ_0 by the vector $\hat{\mu} \in \mathbb{R}^n$ which minimizes the residual squared error $\frac{1}{2}\|Y - \hat{\mu}\|^2$ plus a penalty for each edge $\{i, j\} \in E$ where $\hat{\mu}_i \neq \hat{\mu}_j$. (Here and throughout, $\|\cdot\|$ without a subscript denotes the standard Euclidean norm.) More formally, this estimate $\hat{\mu} \in \mathbb{R}^n$ minimizes the objective function

$$F_0(\mu) := \frac{1}{2}\|Y - \mu\|^2 + \lambda\|D\mu\|_0, \qquad \|D\mu\|_0 := \sum_{\{i,j\}\in E} \mathbb{1}\{\mu_i \neq \mu_j\}. \tag{L0}$$

$D : \mathbb{R}^n \to \mathbb{R}^E$ denotes a vertex-edge incidence matrix with entries in $\{-1, 0, 1\}$ that maps $\mu \in \mathbb{R}^n$ to the vector of edge differences $(\mu_i - \mu_j)_{\{i,j\}\in E}$ (with an arbitrary sign for each edge), and the penalty term $\|D\mu\|_0$ denotes the usual "l_0-norm" of $D\mu$. λ is a tuning parameter that controls the magnitude of this penalty.

For reasons to be discussed, we will also consider procedures that minimize a more general weighted version of the above objective function,

$$F_{\mathrm{w}}(\mu) := \frac{1}{2}\|Y - \mu\|^2 + \lambda\|D\mu\|_{\mathrm{w}}, \qquad \|D\mu\|_{\mathrm{w}} := \sum_{\{i,j\}\in E} \mathrm{w}(i, j)\mathbb{1}\{\mu_i \neq \mu_j\}, \tag{W}$$

where $\mathrm{w} : E \to \mathbb{R}_+$ assigns a non-negative weight to each edge. This allows possibly different penalty values to be applied to different edges of the graph.

The primary purposes of this paper are to advocate for the use of an algorithm by Boykov, Veksler, and Zabih [BVZ01] for minimizing (L0) and (W) for general graphs, to propose minimization of the weighted objective (W) for application to inhomogeneous real-world networks, and to provide a unified theoretical analysis of the squared-error risks of the resulting estimates $\hat{\mu}$ for all graphs. We summarize our main contributions as follows:

(1) While exact minimization of (L0) or (W) may be computationally intractable for general graphs, we show that a polynomial-time algorithm using the α-expansion procedure of [BVZ01] returns an approximate minimizer $\hat{\mu}$ that achieves the same statistical risk guarantees as the exact minimizer (up to constant factors). Computation of $\hat{\mu}$ is reasonably efficient in practice and yields good empirical signal recovery in our tested examples. In this sense, we argue that inference based on minimizing (L0) or (W) is in fact computationally tractable, even for large graphs.

(2) As a corollary of our main theoretical results, we show that for any graph G, the estimate $\hat{\mu}$ (exactly or approximately) minimizing (L0) with $\lambda \asymp \sigma^2 \log|E|$ satisfies an "edge-sparsity" oracle inequality

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \lesssim \inf_{\mu\in\mathbb{R}^n} \|\mu - \mu_0\|^2 + \sigma^2 \max(\|D\mu\|_0, 1) \log|E|. \tag{1}$$

This bounds the squared-error risk of $\hat{\mu}$ in terms of the approximability of μ_0 by a piecewise-constant signal. If it is known that $\|D\mu_0\|_0 \leq s$, then setting instead $\lambda \asymp \sigma^2(1 + \log\frac{|E|}{s})$ yields

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \lesssim \sigma^2 s\left(1 + \log\frac{|E|}{s}\right). \tag{2}$$

The risk bound (2) is rate-optimal in a minimax sense over the edge-sparse signal class $\{\mu_0 : \|D\mu_0\|_0 \leq s\}$ up to a multiplicative factor depending on the mean vertex degree of G.

(3) A popular alternative to minimizing (L0) is to minimize its l_1/total-variation relaxation,

$$F_1(\mu) := \frac{1}{2}\|Y-\mu\|^2 + \lambda\|D\mu\|_1, \qquad \|D\mu\|_1 := \sum_{\{i,j\}\in E} |\mu_i - \mu_j|. \tag{TV}$$

One advantage of this approach is that (TV) is convex and can be exactly minimized in polynomial time. However, we observe that if the underlying signal is truly piecewise constant, then the resulting estimate may be less accurate than approximate minimizers of (L0), particularly in high signal-to-noise regimes. Theoretically, we demonstrate that whether risk guarantees comparable to (1) and (2) hold for $\hat{\mu}$ minimizing (TV) depends on properties of the graph. For example, results of [HR16] imply that (1) is exactly or nearly achieved when G is the 2-D or 3-D lattice, but known risk bounds are much weaker than (1) and (2) when G is the linear chain. For the linear chain, we show that this is not an artifact of the existing analysis, and that in fact

$$\inf_{\lambda\geq 0} \sup_{\mu_0:\|D\mu_0\|_0\leq s} \mathbb{E}[\|\hat{\mu}^\lambda - \mu_0\|^2] \gtrsim \sigma^2(\log n)^{-5}\sqrt{sn},$$

where $\hat{\mu}^\lambda$ minimizes (TV) for each λ. This result is connected to the "slow rate" of convergence in prediction risk for the Lasso [Tib96, CDS01] in certain linear regression settings with correlated predictors.

(4) When G has regions of differing connectivity, we argue that $\|D\mu\|_0$ is not a spatially homogeneous measure of complexity, and it may be more appropriate to minimize the edge-weighted objective (W) where $\mathrm{w}(i,j)$ measures the contribution of edge $\{i,j\}$ to the connectivity of the graph. One such weighting, inspired by the analyses in [SKS13, SSR13], is to weigh each edge by its effective resistance when G is viewed as an electrical resistor network. We observe in simulations of a simple epidemic model on several real-world networks that this weighting can yield a substantial reduction in error over estimators that minimize the unweighted objective (L0). Theoretically, we consider $\hat{\mu}$ (exactly or approximately) minimizing (W) for the more general class of weightings $\mathrm{w} : E \to \mathbb{R}_+$ belonging to the spanning tree polytope of G. We establish that the guarantee (2) holds for this estimate over the larger class $\{\mu_0 : \|D\mu_0\|_{\mathrm{w}} \leq s\}$, and this guarantee is minimax rate-optimal up to a universal (graph-independent) constant factor, for all graphs G.

We provide more detailed discussions of the above points and state our theoretical results in Section 2. Simulation studies comparing minimization of (L0), (W), and (TV) over several graphs are presented in Section 3. Proofs are deferred to the appendices.

1.1. **Related work.** A lot is known about minimization of (L0) in the changepoint setting where G is the linear chain. In this setting, (L0) may be exactly minimized by dynamic programming in quadratic time [AL89, WL02, JSB$^+$05], and pruning ideas have been developed to reduce runtime, which may be near-linear in practice [KFE12]. Correct recovery of changepoint locations and distributional properties of $\hat{\mu}$ minimizing (L0) were studied asymptotically in [Yao88, YA89] when the number of true changepoints s is fixed. Non-asymptotic risk bounds similar to (1) and (2) were established for estimators minimizing similar objectives in [Leb05, BM07]; we discuss this further below. Extension to the recovery of nearly piecewise-constant functions over a continuous interval was considered in [BKL$^+$09].

For other graphs G, with focus on the 2-D lattice for images, many algorithms have been developed to minimize (L0) and similar objectives for discrete Markov random fields, where each μ_i is allowed to take value in a finite set of "labels". One class of such algorithms uses minimum s-t cuts on augmented graphs; we refer the reader to [KZ04] and the contained references for a review. For general graphs and only two distinct labels, [GPS89] showed that the exact (discrete) minimizer may be computed via a single minimum s-t cut. A minimum-cut procedure for binary labels was

applied in [SKS13] to solve a testing/detection analogue of our estimation problem. For G the 2-D lattice with three or more distinct labels, exact minimization of (L0) is NP-hard even in this discrete setting [BVZ01]. Our paper analyzes a graph-cut algorithm from [BVZ01] that applies to more than two labels, where the exact minimization property is lost and replaced by a factor-2 approximation guarantee for the objective function value. One contribution of our paper is to show that the deterministic guarantee of this algorithm may be used to establish a statistical risk bound.

For general graphs G, our characterization of signal complexity by the induced cut-size $\|D\mu_0\|_0$ follows [SKS13, SSR13] which studied similar problems in the testing setting. The testing problem of detecting an anomalous cluster of vertices in a graph was also previously studied in [ACDH05, ACCHZ08, ACCD11, ACG13], with greater focus on achieving optimal detection for specific graph structures and/or specific shapes of the support of μ_0. A connection between this problem and effective edge resistances emerged in the theoretical analyses in [SKS13, SSR13]. In particular, minimization of (W) for $\mathrm{w} : E \to \mathbb{R}_+$ in the spanning tree polytope of G may be viewed as a certain derandomization of the wavelet spanning tree idea of [SSR13] when the underlying tree is chosen at random; we discuss this connection in Remark 2.18.

The estimators $\hat{\mu}$ exactly minimizing (L0) and (W) are examples of general model-complexity penalized estimators studied in [BBM99, BM01, BM07]. In certain cases, the penalties we impose may be smaller than those studied in [BBM99, BM01, BM07]; we discuss this further in Remark 2.17. We control the supremum of a certain Gaussian process using a different argument, which is specialized to the graph-based problem we consider. A theoretical focus of [BBM99, BM01, BM07] was on adaptive attainment of minimax rates over various families of models—for example, for the linear chain graph, [Leb05, BM07] considered penalties increasing but concave in the number of changepoints, and the resulting estimates achieve the minimax guarantee (2) simultaneously for all s. Use of such a penalty is similar to an adaptive choice of λ in (L0), and indeed, [KFE12] proposed to minimize this type of penalized objective by iteratively updating λ in (L0). We will instead apply a different data-driven procedure to choose λ, although we will not study theoretically the adaptivity properties of our approach.

There is a large body of literature regarding minimization of the l_1-relaxation (TV), which is also called total-variation denoising, the fusion estimator, the fused-lasso signal approximator, the edge-lasso, and the graph-fused-lasso. This procedure and generalizations were suggested in different contexts and guises for the linear chain graph in [LF97, MvdG97, DK01, CDS01, TSR$^+$05] and also studied theoretically in [Rin09, HLL10]. For 2-D lattice graphs in image denoising, it was proposed in a slightly different constrained and isotropic form in [ROF92], and later studied also in the anisotropic Lagrangian form as stated in (TV), see e.g. [CL97, GO09]. For more general graphs, this procedure and various generalizations have been studied in [Hoe10, KS11, TT11, SRS12, TS15, HR16, WSST16, SWT16], among others. A body of theoretical work establishes that $\hat{\mu}$ minimizing (TV) is (or is nearly) minimax rate-optimal over signal classes of bounded variation, $\{\mu_0 : \|D\mu_0\|_1 \leq t\}$, for the linear chain graph and higher-dimensional lattices [MvdG97, SWT16, WSST16, HR16]. Several risk bounds for the exact-sparsity classes $\{\mu_0 : \|D\mu_0\|_0 \leq s\}$ that we consider were also established for the linear chain graph in [DHL17, LSRT16] and for general graphs in [HR16, HMPSST16]. We believe that some of the benefits of using effective-resistance-weighted edge penalties may also apply to the l_1/TV setting, although we are not sure if this approach has been previously studied.

1.2. **Acknowledgments.** We would like to thank Professor Amin Saberi for teaching an engaging course on approximation algorithms at Stanford University, which inspired this work. ZF would like to thank Professor James Sharpnack for helpful and inspiring discussions.

1.3. **Notation.** For a vector v, $\|v\| := (\sum_i v_i^2)^{1/2}$ denotes the Euclidean norm, $\|v\|_0 := |\{i : v_i \neq 0\}|$ the "l_0-norm", $\|v\|_1 := \sum_i |v_i|$ the l_1-norm, and $\|v\|_\infty := \max_i |v_i|$ the l_∞-norm. For vectors v and w, $\langle v, w\rangle = \sum_i v_i w_i$ denotes the Euclidean inner-product.

$\mathbb{R}_+$ denotes the non-negative reals. For an edge weighting $\mathrm{w} : E \to \mathbb{R}_+$, $\mathrm{w}(i,j)$ is shorthand for $\mathrm{w}(\{i,j\})$, and we denote $\mathrm{w}(E') := \sum_{\{i,j\}\in E'} \mathrm{w}(i,j)$ for any edge subset $E' \subseteq E$. For two edge weightings $\mathrm{w}, \mathrm{r} : E \to \mathbb{R}_+$, we write $\mathrm{w} \geq \mathrm{r}$ if $\mathrm{w}(i,j) \geq \mathrm{r}(i,j)$ for all edges $\{i,j\} \in E$. For $v \in \mathbb{R}^E$, $\|v\|_{\mathrm{w}} := \sum_{\{i,j\}\in E} \mathrm{w}(i,j)v_{\{i,j\}}$ denotes the l_0-norm weighted by w.

For positive quantities a and b, we write $a \lesssim b$ if $a \leq Cb$ and $a \asymp b$ if $ca \leq b \leq Ca$, where $C, c > 0$ are universal constants (independent of the graph G).

$\mathbb{1}\{\cdot\}$ denotes the indicator function, i.e. $\mathbb{1}\{\mathcal{E}\} = 1$ if condition $\mathcal{E}$ is true and 0 otherwise.

2. Algorithm and theoretical results

2.1. **Approximation algorithm.** As discussed in Section 1.1, whether (L0) and (W) may be minimized exactly in polynomial time depends on the graph G. However, good approximation of the solution is tractable for any graph. We review in this section one approach that achieves such an approximation, based on discretizing the range of values of the entries of μ and applying the α-expansion local move of [BVZ01] for discrete Markov random fields. We will describe the algorithm for minimizing (W), as (L0) is clearly a special case.

The fundamental property of this algorithm will be that its output is a $(\tau, \delta\mathbb{Z})$-local-minimizer for the objective function (W), defined as follows:

Definition 2.1. For $\delta > 0$, denote by

$$\delta\mathbb{Z} := \{\ldots, -3\delta, -2\delta, -\delta, 0, \delta, 2\delta, 3\delta, \ldots\}$$

the set of all integer multiples of δ. For any $\mu \in \mathbb{R}^n$, a **$\delta\mathbb{Z}$-expansion** of μ is any other vector $\tilde{\mu} \in \mathbb{R}^n$ such that there exists a single value $c \in \delta\mathbb{Z}$ for which, for every $i \in \{1, \ldots, n\}$, either $\tilde{\mu}_i = \mu_i$ or $\tilde{\mu}_i = c$. For $\delta > 0$ and $\tau \geq 0$, a **$(\tau, \delta\mathbb{Z})$-local-minimizer** of (W) is any $\mu \in \mathbb{R}^n$ such that for every $\delta\mathbb{Z}$-expansion $\tilde{\mu}$ of μ,

$$F_{\mathrm{w}}(\mu) - \tau \leq F_{\mathrm{w}}(\tilde{\mu}).$$

More informally, a $\delta\mathbb{Z}$-expansion of μ can replace any subset of vertex values by a single new value, and a $(\tau, \delta\mathbb{Z})$-local-minimum is such that no further $\delta\mathbb{Z}$-expansion reduces the objective value by more than τ.[1] This definition does not require $(\tau, \delta\mathbb{Z})$-local-minima to necessarily have all of its entries belonging to $\delta\mathbb{Z}$—hence, in particular, any global minimizer of (W) is also a $(\tau, \delta\mathbb{Z})$-local-minimizer for any $\delta > 0$ and $\tau \geq 0$. We define analogously $(\tau, \delta\mathbb{Z})$-local-minimizers for (L0).

The α-expansion procedure of [BVZ01] may be used to compute a $(\tau, \delta\mathbb{Z})$-local-minimizer efficiently with graph cuts. We review the details of this procedure and how we will apply it to our problem in Algorithm 1. The quantity $\bar{Y}^\delta$ in Line 1 denotes the overall sample mean rounded to the nearest value in $\delta\mathbb{Z}$, and $Y_{\min}$ and $Y_{\max}$ in Line 3 denote the minimum and maximum values of Y_i. We note that sometimes graph cut procedures are applied in computer vision in a "classification" context with a small number of labels so as to yield a coarse segmentation/clustering of the image. This is not our focus, and we will instead apply and study Algorithm 1 under a very fine discretization δ so as to yield a good solution to the original continuous problem.

We may formally establish that this algorithm runs in polynomial time; we discuss empirical runtime in Section 3.2.

Proposition 2.2. *Suppose $Y_{\max} - Y_{\min}$, $1/\delta$, and $1/\tau$ are at most polynomial in n. Then Algorithm 1 always terminates in time polynomial in n.*

Proof. Each iteration of the loop in Line 3 is computable in time polynomial in n. The number of times this loop is executed in each outer loop of Line 2 is polynomial in n, by the conditions on $Y_{\max} - Y_{\min}$ and $1/\delta$. Every time $\hat{\mu}$ is updated in Line 18, the objective value $F_{\mathrm{w}}(\hat{\mu})$ must decrease

[1] We introduce a slack $\tau \geq 0$ for theoretical purposes, as a simple device to provably ensure that the algorithm terminates in polynomial time. In practice, we find that the algorithm has reasonable runtime even with $\tau = 0$.

Algorithm 1 α-expansion algorithm [BVZ01] to find a $(\tau, \delta\mathbb{Z})$-local-minimizer of (W)

1: Initialize $\hat{\mu} \in \mathbb{R}^n$ by setting $\hat{\mu}_i = \bar{Y}^\delta$ for all $i \in \{1, \dots, n\}$.
2: **loop**
3: **for** each $c \in \delta\mathbb{Z} \cap [Y_{\min} - \delta, Y_{\max} + \delta]$ **do**
4: Construct the following edge-weighted augmentation $G_{c,\hat{\mu}}$ of the original graph G:
5: Introduce a source vertex s and a sink vertex t.
6: Connect s to each $i \in \{1, \dots, n\}$ with weight $\frac{1}{2}(Y_i - c)^2$.
7: Connect t to each $i \in \{1, \dots, n\}$ with weight $\frac{1}{2}(Y_i - \hat{\mu}_i)^2$ if $\hat{\mu}_i \neq c$, or weight ∞ if $\hat{\mu}_i = c$.
8: **for** each edge $\{i, j\} \in E$ **do**
9: **if** $\hat{\mu}_i = \hat{\mu}_j$ **then**
10: Assign weight $\lambda \mathrm{w}(i, j) \mathbb{1}\{\hat{\mu}_i \neq c\}$ to $\{i, j\}$.
11: **else**
12: Introduce a new vertex $\mathsf{a}_{i,j}$.
13: Replace $\{i, j\}$ by the three edges $\{i, \mathsf{a}_{i,j}\}$, $\{j, \mathsf{a}_{i,j}\}$, and $\{\mathsf{t}, \mathsf{a}_{i,j}\}$, with weights $\lambda \mathrm{w}(i, j) \mathbb{1}\{\hat{\mu}_i \neq c\}$, $\lambda \mathrm{w}(i, j) \mathbb{1}\{\hat{\mu}_j \neq c\}$, and $\lambda \mathrm{w}(i, j)$, respectively.
14: **end if**
15: **end for**
16: Find the minimum s-t cut (S, T) of $G_{c,\hat{\mu}}$ such that $\mathsf{s} \in S$, $\mathsf{t} \in T$.
17: **if** the value of the minimum cut is less than $F_{\mathrm{w}}(\hat{\mu}) - \tau$ **then**
18: For each vertex $i \in \{1, \dots, n\}$ belonging to T, set $\hat{\mu}_i = c$.
19: **end if**
20: **end for**
21: **if** no vertex value $\hat{\mu}_i$ was changed **then**
22: **return** $\hat{\mu}$
23: **end if**
24: **end loop**

by at least τ. The initial objective value $\frac{1}{2}\|Y - \bar{Y}^\delta\|^2$ is polynomial in n, and hence the number of iterations of the outer loop is also polynomial in n by the condition on $1/\tau$. □

The following verifies that Algorithm 1 indeed returns a local minimizer of (W) in the sense of Definition 2.1:

Proposition 2.3. *The vector $\hat{\mu}$ returned by Algorithm 1 is a $(\tau, \delta\mathbb{Z})$-local-minimizer of (W).*

Proof. The objective (W) may be written as

$$F_{\mathrm{w}}(\mu) = \frac{1}{2}\sum_{i=1}^{n}(Y_i - \mu_i)^2 + \sum_{\{i,j\}\in E} \lambda \mathrm{w}(i, j) \mathbb{1}\{\mu_i \neq \mu_j\},$$

and each edge-cost $\lambda \mathrm{w}(i, j) \mathbb{1}\{\mu_i \neq \mu_j\}$ satisfies the triangle inequality in the sense

$$\lambda \mathrm{w}(i, j) \mathbb{1}\{a \neq c\} \leq \lambda \mathrm{w}(i, j) \mathbb{1}\{a \neq b\} + \lambda \mathrm{w}(i, j) \mathbb{1}\{b \neq c\}$$

for any $a, b, c \in \mathbb{R}$. Hence, Theorem 5.4 and Corollary 5.5 of [BVZ01] guarantee the following property for the minimum s-t cut (S, T) computed in Line 16 of Algorithm 1:

(*) Let $\mathcal{V}$ denote the set of all $\delta\mathbb{Z}$-expansions of $\hat{\mu}$ where the new vertex value is $c \in \delta\mathbb{Z}$. Then $\operatorname{argmin}_{\mu\in\mathcal{V}} F_{\mathrm{w}}(\mu)$ is precisely the vector $\tilde{\mu}$ such that $\tilde{\mu}_i = \hat{\mu}_i$ if $i \in S$ and $\tilde{\mu}_i = c$ if $i \in T$, and the value of the minimum cut is equal to $F_{\mathrm{w}}(\tilde{\mu})$.

If the returned vector $\hat{\mu}$ is not a $(\tau, \delta\mathbb{Z})$-local-minimizer, then there exists a $\delta\mathbb{Z}$-expansion $\tilde{\mu}$ of $\hat{\mu}$, changing some values of $\hat{\mu}_i$ to $c \in \delta\mathbb{Z}$, that decreases F_{w} by more than τ. Then there must exist

such a $\delta\mathbb{Z}$-expansion $\tilde{\mu}$ with $c \in [Y_{\min}-\delta, Y_{\max}+\delta]$, as the entries of $\hat{\mu}$ and Y all belong to this range. But the above property (*) ensures that the value of the minimum cut in Line 17 computed for this value of c is at most $F_{\mathrm{w}}(\tilde{\mu})$, and hence termination of the algorithm implies $F_{\mathrm{w}}(\hat{\mu}) - \tau < F_{\mathrm{w}}(\tilde{\mu})$, contradicting that $\tilde{\mu}$ decreases F_{w} by more than τ. □

The l_2 vertex-cost $(Y_i - \mu_i)^2$ and l_0 edge cost $\mathbb{1}\{\mu_i \neq \mu_j\}$ of (L0) and (W) are not intrinsic to this algorithm, and the same method may be applied to approximately minimize

$$F(\mu) = \sum_{i=1}^{n} c_i(Y_i, \mu_i) + \sum_{\{i,j\}\in E} c_{i,j}(\mu_i, \mu_j)$$

for any vertex cost functions c_i and edge cost functions $c_{i,j}$, where each $c_{i,j}$ satisfies a triangle inequality. Thus the algorithm is easily applied to more general data types and likelihood models (although discretization of the parameter space becomes problematic if its dimensionality increases). An l_0 term $\mathbb{1}\{\mu_i \neq 0\}$ may be added to the vertex cost if it is believed that the true signal is both piecewise-constant and sparse, and the edge cost may be replaced by a combination of l_0, l_1, and l_2 penalties. We will not explore these choices in this paper.

2.2. **Theoretical guarantees for l_0 denoising.** In this section, we describe squared-error risk guarantees for $\hat{\mu}$ (exactly or approximately) minimizing the unweighted objective function (L0). Although these results are corollaries of those in Section 2.4 for the weighted objective (W), we state them here separately as they are simpler to understand and also provide a basis for comparison with l_1/total-variation denoising discussed in the next section. We recall that G is fully connected, and we assume throughout that $n \geq 3$.

Recall Definition 2.1 of $(\tau, \delta\mathbb{Z})$-local-minimizers of (L0), which include both the exact minimizer and the estimator computed by Algorithm 1. A sparsity-oracle inequality for any such minimizer holds when the penalty λ in (L0) is set to a "universal" level $C\sigma^2 \log|E|$:

Theorem 2.4. *Let $\delta \leq \sigma/\sqrt{n}$ and $\tau \leq \sigma^2$. For any $\eta > 0$, there exist constants $C_\eta, C'_\eta > 0$ depending only on η such that if $\lambda \geq C_\eta \sigma^2 \log|E|$ and $\hat{\mu}$ is any $(\tau, \delta\mathbb{Z})$-local-minimizer of (L0), then*

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \leq \inf_{\mu\in\mathbb{R}^n} (1+\eta)\|\mu - \mu_0\|^2 + C'_\eta \lambda \max(\|D\mu\|_0, 1). \tag{3}$$

The upper bound in (3) trades off the edge-sparsity of μ and its approximation of the true signal μ_0. Setting $\lambda = C\sigma^2 \log|E|$ yields the guarantee (1) described in the introduction. If the true signal μ_0 has edge-sparsity $\|D\mu_0\|_0 = s$, then evaluating (3) at $\mu = \mu_0$ yields a risk bound of order $\sigma^2 s \log|E|$. When s is known, we may obtain the tighter guarantee (2) by using a smaller penalty:

Theorem 2.5. *Let $\delta \leq \sigma/\sqrt{n}$ and $\tau \leq \sigma^2$. There exist universal constants $C, C' > 0$ such that for any $s \in [1, |E|]$, if $\lambda \geq C\sigma^2(1 + \log\frac{|E|}{s})$ and $\hat{\mu}$ is any $(\tau, \delta\mathbb{Z})$-local-minimizer of (L0), then*

$$\sup_{\mu_0 : \|D\mu_0\|_0 \leq s} \mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \leq C'\lambda s. \tag{4}$$

Theorems 2.4 and 2.5 are reminiscent of estimation guarantees in the sparse normal-means problem: For estimating a signal $\mu_0 \in \mathbb{R}^n$ with at most k nonzero entries, in an asymptotic setting where $k/n \to 0$, the minimax squared-error risk is

$$\inf_{\hat{\mu}} \sup_{\mu_0 : \|\mu_0\| \leq k} \mathbb{E}[\|\hat{\mu} - Y\|^2] \sim 2\sigma^2 k \log\frac{n}{k} \tag{5}$$

and is achieved by the estimator $\hat{\mu} = \mathrm{argmin}_\mu \frac{1}{2}\|Y - \mu\|^2 + \lambda\|\mu\|_0$ for $\lambda = \sigma^2 \log\frac{n}{k}$, corresponding to entrywise hard-thresholding at level $\sqrt{2\lambda}$ [Joh15, Theorem 8.20]. Setting $\lambda = \sigma^2 \log n$ hard-thresholds instead at the universal level $\sqrt{2\sigma^2 \log n}$, and Lemma 1 of [DJ94] implies an oracle

bound

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \leq \inf_{\mu \in \mathbb{R}^n} 1.2\|\mu - \mu_0\|^2 + \sigma^2(2\log n + 1)(\|\mu\|_0 + 1)$$

for any true signal $\mu_0 \in \mathbb{R}^n$. (The infimum is attained at a vector μ such that each entry μ_i equals 0 or $\mu_{0,i}$. The same bound holds with leading coefficient 1 instead of 1.2 if soft-thresholding is applied instead at the same universal level.)

When there is an underlying graph G, the entrywise sparsity condition $\|\mu_0\|_0 \leq k$ may be considered a notion of "vertex" sparsity, in contrast to our notion of edge-sparsity. If G has maximum vertex degree d, then $\{\mu_0 : \|\mu_0\|_0 \leq s/d\} \subseteq \{\mu_0 : \|D\mu_0\|_0 \leq s\}$. (A converse containment is, of course, is not true—there are many signals μ_0 that are piecewise-constant on G, but such that $\|\mu_0\|_0$ is not small.) When d is a small constant, for instance in the cases of the linear chain and 2-D/3-D lattice graphs, Theorem 2.5 implies that the minimax risk over the edge-sparse class $\{\mu_0 : \|D\mu_0\|_0 \leq s\}$ is of the same order as that over the contained vertex-sparse sub-class $\{\mu_0 : \|\mu_0\|_0 \leq s/d\}$.

This type of containment also implies the following minimax lower-bound:

Theorem 2.6. *Suppose G has average vertex degree d. There exists a universal constant $c > 0$ such that for any $s \in [4d, |E|]$,*

$$\inf_{\hat{\mu}} \sup_{\mu_0 : \|D\mu_0\|_0 \leq s} \mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \geq c\sigma^2 \frac{s}{d}\left(1 + \log\frac{|E|}{s}\right), \tag{6}$$

where the infimum is taken over all possible estimators $\hat{\mu} := \hat{\mu}(Y)$.

When the average vertex degree d is not small, there is a gap between the upper bound (4) and lower bound (6) of order d, which we will discuss in Section 2.4.

Remark 2.7. We assume $s \geq 4d$ for the lower bound (6) so that the result does not depend on the exact structure of near-minimum cuts in G. For example, suppose that vertices $\{1, \ldots, n-1\}$ of G are connected in a single cycle with $n-1$ edges, and in addition vertex n is connected to vertex 1 by a single edge. For $s = 1$, any μ_0 with $\|D\mu_0\|_0 \leq 1$ must take a constant value over vertices $\{1, \ldots, n-1\}$ and a possibly different value on vertex n. The minimax squared-error risk over this class is then $2\sigma^2$, rather than being of order $\sigma^2 \log n$. On the other hand, this phenomenon does not occur for the linear chain, where (6) holds for all $s \in [1, |E|]$. Hence some restriction to $s \geq c$ for a small constant c is required for a general lower bound that is independent of the detailed structure of G. Considering the tensor products of these two examples with the complete graph on d vertices, a similar argument shows that in fact a general lower bound must restrict to $s \geq cd$.

While the main focus of our work is estimation, we conclude this section with a result that is of relevance to testing:

Theorem 2.8. *Let $\delta \leq \sigma/\sqrt{n}$ and $\tau \leq \sigma^2$, and suppose μ_0 is constant over G. There exist universal constants $C, C' > 0$ such that if $\lambda \geq C\sigma^2 \log|E|$ and $\hat{\mu}$ is any $(\tau, \delta\mathbb{Z})$-local-minimizer of (L0), then*

$$\mathbb{P}[\hat{\mu} \text{ is constant over } G] \geq 1 - C'n^{-3}.$$

This implies that we may test the null hypothesis

$$H_0 : \mu_0 \text{ is constant} \tag{7}$$

by setting $\lambda \asymp \sigma^2 \log|E|$ and rejecting H_0 if $\hat{\mu}$ is not constant. Denoting by $P^\perp$ the orthogonal projection onto the space orthogonal to the all-1's vector, since

$$\min_{\mu : \mu \text{ is constant}} \|\mu - \mu_0\|^2 = \|P^\perp \mu_0\|^2,$$

the risk bound (4) (or more precisely, Lemma A.3(b) which establishes an analogous bound in probability) implies that this test can distinguish a non-constant alternative μ_0 with high probability as long as $\|P^\perp\mu_0\|^2 \gtrsim \sigma^2\|D\mu_0\|_0 \log|E|$.

2.3. **Comparison with l_1/total-variation denoising.** We compare the guarantees of the preceding section with those attainable by $\hat{\mu}$ minimizing (TV). We focus our discussion on risk guarantees for exactly piecewise-constant signals without any restriction on $\|D\mu_0\|_1$, for more direct comparison with our preceding results. We believe that at present, there is only a partial understanding of the squared-error risk of the TV-estimator $\hat{\mu}$ for general graphs G. This question is related to understanding the prediction risk achievable by the Lasso estimate for sparse linear regression when the design matrix may have poor restricted eigenvalue/compatibility constants, which is a subject of recent research [vdGL13, HL13, DHL17].

One general result for the risk of $\hat{\mu}$ was established in [HR16]. For an arbitrary graph G, let d denote its mean vertex degree, $D : \mathbb{R}^n \to \mathbb{R}^E$ its vertex-edge incidence matrix, $S = D^\dagger$ the Moore-Penrose pseudo-inverse of D, and ρ the maximum Euclidean norm of any column of S. Theorem 2 of [HR16] implies,[2] for the estimator $\hat{\mu}$ minimizing (TV) with the choice $\lambda = \sigma\rho\sqrt{2(1+\log(|E|/\delta))}$, and for any $\mu \in \mathbb{R}^n$, with probability at least $1-2\delta$,

$$\|\hat{\mu}-\mu_0\|^2 \le \|\mu-\mu_0\|^2 + 8\sigma^2\left(\frac{\rho^2\|D\mu\|_0}{\kappa^2}\left(1+\log\frac{|E|}{\delta}\right)+\log\frac{e}{\delta}\right), \tag{8}$$

where κ is a compatibility constant bounded as $\kappa^{-2} \le 4\min(d, \|D\mu\|_0)$.

An important difference between this result and Theorem 2.4 of the preceding section lies in the appearance of ρ^2, which is graph-dependent. Assuming G has small average degree d, the above guarantee is similar to Theorem 2.4 if ρ^2 is small. It is shown in [HR16] that $\rho^2 \lesssim 1$ for 3-D (and higher-dimensional) lattice graphs and $\rho^2 \lesssim \log n$ for 2-D lattice graphs, indicating that the total-variation estimator $\hat{\mu}$ is nearly minimax rate-optimal over $\{\mu_0 : \|D\mu_0\|_0 \le s\}$ in these settings. However, when G is, for example, the linear chain, $\rho^2 \asymp n$ and the guarantee (8) is larger than those established in the preceding section by this factor of n.

More specialized analyses were performed for the linear chain in [DHL17, LSRT16]. To state one result from the latter work, for G the linear chain, let $\Delta(\mu_0)$ denote the minimum distance between two changepoints of μ_0. More precisely, denoting by $1 \le i_1 < \ldots < i_s < n$ the values i for which $\mu_{0,i} \ne \mu_{0,i+1}$ and letting $i_0 := 0$ and $i_{s+1} = n$, define $\Delta(\mu_0) := \min_{0\le r\le s} i_{r+1} - i_r$. Then Theorem 4 of [LSRT16] shows, if $\|D\mu_0\|_0 = s$ and $\lambda = \sigma(n\Delta(\mu_0))^{1/4}$, then

$$\mathbb{E}[\|\hat{\mu}-\mu_0\|^2] \lesssim \sigma^2 s\left((\log s + \log\log n)\log n + \sqrt{n/\Delta(\mu_0)}\right).$$

When $\Delta(\mu_0) \lesssim n^\alpha$ for any $\alpha < 1$, this bound differs from the guarantee of Theorem 2.5 by a factor of roughly $n^{(1-\alpha)/2}$. A similar result was recently obtained in [HMPSST16] for general graphs, where $\Delta(\mu_0)$ measures the minimum spacing between changepoints in an associated linear chain graph induced by a depth-first-search traversal of G.

To clarify that the suboptimality of $\hat{\mu}$ indicated by these results is not an artifact of the analysis, we state in this section two lower bounds for the squared-error risk of total-variation denoising for the linear chain; proofs are given in Appendix B. First, when the true signal μ_0 is zero, we have the following behavior for $\hat{\mu}$:

Theorem 2.9. *Let G be the linear chain graph with n vertices and $n-1$ edges, and suppose $\mu_0 = 0$. There exists a constant $c > 0$ such that for any fixed $\lambda \in [0, \sigma\sqrt{cn/\log n}]$, if $\hat{\mu}$ is the minimizer of (TV), then the following hold:*

(a) For some constants $C, c' > 0$, letting $\hat{k} := \|D\hat{\mu}\|_0 + 1$ be the number of constant intervals of $\hat{\mu}$,

$$\mathbb{P}\left[\hat{k} > \frac{cn}{\max(\lambda^2/\sigma^2, 1)\log n}\right] \ge 1 - Ce^{-c'n/\max(\lambda^2/\sigma^2,1)}.$$

[2]The result of [HR16] is more general and stated in terms of a certain trade-off between $\|D\mu\|_1$ and $\|D\mu\|_0$; for purposes of discussion, we have specialized to the "pure-l_0" setting.

(b) For some constant $c' > 0$, the squared-error risk of $\hat{\mu}$ satisfies

$$\mathbb{E}[\|\hat{\mu}\|^2] \geq \frac{c'\sigma^2 n}{\max(\lambda^2/\sigma^2, 1)(\log n)^4}.$$

Hence if $\mu_0 = 0$ and $\lambda \asymp \sigma n^\alpha$ for any $\alpha < 1/2$, then the number of changepoints and the squared-error risk of the TV-estimator $\hat{\mu}$ are (up to logarithmic factors) at least of order $n^{1-2\alpha}$ and $\sigma^2 n^{1-2\alpha}$, respectively. As a consequence, we obtain the following lower bound in a minimax sense:

Theorem 2.10. *Let G be the linear chain graph with n vertices and $n-1$ edges, and let $\Delta(\mu_0)$ denote the minimum distance between two changepoints in μ_0. For each fixed $\lambda \geq 0$, let $\hat{\mu}^\lambda$ denote the minimizer of (TV) for this λ. Then there exists a constant $c > 0$ such that for any $s \in [2, n-1]$ and $\Delta \leq n/(s+1)$,*

$$\inf_{\lambda \geq 0} \sup_{\mu_0 : \|D\mu_0\|_0 \leq s,\, \Delta(\mu_0) \geq \Delta} \mathbb{E}[\|\hat{\mu}^\lambda - \mu_0\|^2] \geq \frac{c\sigma^2}{(\log n)^5} \sqrt{\frac{sn}{\Delta}}.$$

In particular, setting $\Delta = 1$ removes restrictions on the minimum spacing between changepoints and yields the lower-bound $c\sigma^2\sqrt{sn}/(\log n)^5$ over the class $\{\mu_0 : \|D\mu_0\|_0 \leq s\}$. Let us remark that this does not contradict the upper bound $\|\hat{\mu} - \mu_0\|^2 \lesssim \sigma^2 n^{1/3} t^{2/3}$ over the class $\{\mu_0 : \|D\mu_0\|_1 \leq t\}$ established in [MvdG97], as our proof of the above result considers for each λ a signal μ_0 with maximum value λ/Δ, so that $\|D\mu_0\|_1$ is larger than $\|D\mu_0\|_0$ by this factor. This is in agreement with our empirical observations in Section 3 that the estimate minimizing (L0) improves over that minimizing (TV) particularly in high signal-to-noise regimes.

Theorem 2.10 may be re-interpreted in the context of the Lasso estimate for sparse linear regression: Setting $\beta_0 := D\mu_0$ and

$$X := D^\dagger = \begin{pmatrix} -\frac{n-1}{n} & -\frac{n-2}{n} & -\frac{n-3}{n} & \cdots & -\frac{1}{n} \\ \frac{1}{n} & -\frac{n-2}{n} & -\frac{n-3}{n} & \cdots & -\frac{1}{n} \\ \frac{1}{n} & \frac{2}{n} & -\frac{n-3}{n} & \cdots & -\frac{1}{n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{1}{n} & \frac{2}{n} & \frac{3}{n} & \cdots & \frac{n-1}{n} \end{pmatrix} \in \mathbb{R}^{n \times (n-1)},$$

minimizing (TV) is equivalent to minimizing the Lasso objective

$$\frac{1}{2}\|\tilde{Y} - X\beta\|^2 + \lambda\|\beta\|_1$$

over $\beta \in \mathbb{R}^{n-1}$, where $\tilde{Y} = (Y_1 - \bar{Y}, \ldots, Y_n - \bar{Y})$ denotes Y centered by its mean. The maximum column norm of X is $\sqrt{n/4}$, the error $\|\hat{\mu} - \mu_0\|^2$ corresponds to n times the "prediction loss" $n^{-1}\|X\hat{\beta} - X\beta_0\|^2$, and in this context Theorem 2.10 implies

$$\inf_{\lambda \geq 0} \sup_{\beta_0 : \|\beta_0\|_0 \leq s} \mathbb{E}\left[\frac{1}{n}\|X\hat{\beta}^\lambda - X\beta_0\|^2\right] \geq \frac{c\sigma^2}{(\log n)^5}\sqrt{\frac{s}{n}}.$$

Hence the minimax prediction risk for the Lasso estimate $\hat{\beta}$ over $\{\beta_0 : \|\beta_0\|_0 \leq s\}$ decays essentially no faster than order $n^{-1/2}$. This is in contrast to the faster rate of n^{-1} that is achievable when X has well-behaved restricted eigenvalue constants (see, for example, [BRT09, VDGB09] and the references contained therein).

More generally, for any connected graph G, noting that $D^\dagger \in \mathbb{R}^{n \times E}$ is of rank $n-1$ with range orthogonal to the all-1's vector, minimizing (TV) is equivalent to minimizing

$$\frac{1}{2}\|\tilde{Y} - D^\dagger\beta\|^2 + \lambda\|\beta\|_1 \text{ subject to } \beta \in \text{range}(D),$$

where range(D) denotes the column span of D in $\mathbb{R}^E$. The results of the two preceding sections imply that whenever G has small average degree, the "fast" optimal rate of convergence for prediction

risk over the class $\{\beta_0 \in \text{range}(D) : \|\beta_0\|_0 \leq s\}$ for the above problem is attainable in polynomial time,[3] even if it is not achieved by the l_1 estimator. This may be contrasted with the negative results of [ZWJ14, ZWJ15], which show that there exist adversarial design matrices X for sparse regression (not based on any graph structure) where such a fast rate is not achieved by a broad class of M-estimators or by any polynomial-time algorithm returning an s-sparse output.

2.4. **Edge-weighting for inhomogeneous graphs.** In this section, we generalize the results of Section 2.2 by considering (exact or approximate) minimizers $\hat{\mu}$ of the edge-weighted objective (W). We again assume implicitly that G is fully connected and $n \geq 3$.

Let us motivate our discussion by the following example, which examines the factor-d gap between the minimax upper and lower bounds of (4) and (6):

Example 2.11. Let G be the complete graph on n vertices. Then the average vertex degree of G is $d = n - 1$, and (6) implies

$$\inf_{\hat{\mu}} \sup_{\mu_0 : \|D\mu_0\|_0 \leq s} \mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \gtrsim \sigma^2 \frac{s}{n} \left(1 + \log \frac{|E|}{s}\right). \tag{9}$$

This lower bound is in fact tight, and the upper bound of (4) is loose by a factor of n: Theorem 2.15 below will imply that setting $\lambda \asymp \frac{\sigma^2}{n}(1 + \log \frac{|E|}{s})$ in (L0) achieves the above level of risk, when G is the complete graph.

On the other hand, let G be a "tadpole" graph consisting of a linear chain of $n/2$ vertices with one endpoint connected by an edge to a clique of $n/2$ remaining vertices. The average vertex degree of G in this case is $d = (n+1)/2$, so a direct application of (6) still yields (9). However, by restricting to the sub-class of signals that take a constant value on the $n/2$-clique, it is clear that the minimax risk over $\{\mu_0 : \|D\mu_0\|_0 \leq s\}$ is at least that of estimating the signal over only the linear chain portion of G with $n/2$ vertices. The lower bound (6) applied to only this sub-graph implies that in this case, the upper bound (4) is tight up to a constant factor, and the lower bound (9) is loose by a factor of $1/n$.

This example highlights the problem that the complexity measure $\|D\mu_0\|_0$ is not necessarily spatially homogeneous over G. For example, when G is the tadpole graph, a signal μ_0 that is constant over all but one vertex belonging to the $n/2$-clique has $\|D\mu_0\|_0 = n/2 - 1$, but a signal taking a different value at each of the vertices of the linear chain also has $\|D\mu_0\|_0 = n/2 - 1$. A theoretical consequence is that the minimax risk over $\{\mu_0 : \|D\mu_0\|_0 \leq s\}$ is controlled by the least well-connected portion of the graph. A practical consequence is that any choice of λ will either oversmooth the signal over the $n/2$-clique or undersmooth the signal over the linear chain, and no single choice of λ leads to good signal recovery in both of these regions.

While the tadpole graph is an extreme example, the same phenomenon arises in any graph with regions of varying connectivity. In such applications, we propose to consider the weighted objective function (W) where each edge is weighted by a measure of its contribution to the connectivity of G. We believe both that minimizing this weighted objective is usually a more reasonable procedure in practice and that the value $\|D\mu_0\|_{\mathrm{w}}$ provides a better indication of the complexity of the piecewise-constant signal μ_0.

One specific weighting $\mathrm{w} : E \to \mathbb{R}_+$ that implements this idea is to weigh each edge by its effective resistance.

Definition 2.12. Let G be a connected graph and $\{i, j\}$ an edge in G. The **effective resistance** $\mathrm{r}(i, j)$ of this edge has the following four equivalent definitions:

(1) $\mathrm{r}(i, j)$ is the effective electrical resistance measured across vertices i and j when G represents an electrical network where each edge is a resistor with resistance 1.

[3]For polynomial runtime of Algorithm 1, we require $\|\beta_0\|_\infty$ to be polynomial in n. The bound for $\|\beta_0\|_\infty$ influences only the runtime and not the statistical risk guarantee.

(2) Let L be the (unweighted) Laplacian matrix of G, $L^\dagger$ the pseudo-inverse of L, and e_i the basis vector with i^{th} entry 1 and remaining entries 0. Then $\mathrm{r}(i,j) = (e_i - e_j)L^\dagger(e_i - e_j)$.
(3) Consider a random walk on G starting at vertex i, and let t be the number of steps taken to reach vertex j and then return to vertex i for the first time. Then $\mathrm{r}(i,j) = \frac{1}{2|E|}\mathbb{E}[t]$.
(4) Let T be (the edges of) a random spanning tree of G chosen uniformly from the set of all spanning trees of G. Then $\mathrm{r}(i,j) = \mathbb{P}[\{i,j\} \in T]$.

For verification of the equivalence of the above definitions, we refer the reader to the survey [L93] and the introductory sections of [GBS08]. In practice, $\mathrm{r}(i,j)$ may be computed via the second characterization above using fast Laplacian solvers (e.g. those in [ST04, KMP10, SS11, LB12]).

The fourth characterization above describes one sense in which $\mathrm{r}(i,j)$ measures the "contribution" of edge $\{i,j\}$ to the connectivity of G: For example, if removing $\{i,j\}$ breaks G into two disconnected components, then every spanning tree T of G must contain $\{i,j\}$, so $\mathrm{r}(i,j) = 1$. Conversely, if there are many short alternative paths from i to j not using edge $\{i,j\}$, then $\mathrm{r}(i,j)$ is much smaller than 1.

More generally, a definition of the contribution of each edge to the graph connectivity is given by any weighting belonging to the spanning tree polytope of G:

Definition 2.13. A weighting $\mathrm{r} : E \to \mathbb{R}_+$ is a *spanning tree weighting* if there exists a spanning tree T of G such that $\mathrm{r}(i,j) = 1$ if $\{i,j\} \in T$ and $\mathrm{r}(i,j) = 0$ otherwise. The **spanning tree polytope** $\mathcal{ST}(G)$ is the convex hull of all spanning tree weightings.

Thus, if $\mathrm{r} \in \mathcal{ST}(G)$, then there exist spanning trees $T_1, \dots, T_M$ of G and $\lambda_1, \dots, \lambda_M > 0$ with $\sum_{m=1}^M \lambda_m = 1$ such that

$$\mathrm{r}(i,j) = \sum_{m=1}^M \lambda_m \mathbb{1}\{\{i,j\} \in T_m\}$$

for every edge $\{i,j\} \in E$. The weighting r is thus identified with the probability distribution of a random spanning tree T, where $T = T_m$ with probability λ_m. This distribution satisfies the property, for all $\{i,j\} \in E$,

$$\mathrm{r}(i,j) = \mathbb{P}[\{i,j\} \in T].$$

For any subset of edges $E' \subseteq E$ and weighting $\mathrm{w} : E \to \mathbb{R}_+$, let us denote $\mathrm{w}(E') = \sum_{\{i,j\} \in E'} \mathrm{w}(i,j)$ as the total weight of these edges. Then the above implies, for the random spanning tree T associated to $\mathrm{r} \in \mathcal{ST}(G)$,

$$\mathrm{r}(E') = \mathbb{E}[|E' \cap T|]. \tag{10}$$

The effective resistance weighting of Definition 2.12 corresponds to the uniform distribution for T.

The results below describe the squared-error risk of the estimator $\hat{\mu}$ that (exactly or approximately) minimizes (W) for any edge-weighting $\mathrm{w} : E \to \mathbb{R}_+$. We derive, for all graphs G, minimax upper and lower bounds on this risk over the class $\{\mu_0 : \|D\mu_0\|_{\mathrm{w}} \le s\}$. The tightness of these bounds will depend on how close w is to the spanning tree polytope $\mathcal{ST}(G)$—for effective resistance weighting, or more generally for any $\mathrm{r} \in \mathcal{ST}(G)$, these bounds are tight up to a universal constant factor independent of the graph.

Let us make a remark about scaling, which is important for the interpretation of the below results: As rescaling w by $c > 0$ and λ by $1/c$ leads to the same penalty in (W), we will state all of our results, for simplicity and without loss of generality,[4] under a scaling such that $\mathrm{w} \ge \mathrm{r}$ for some $\mathrm{r} \in \mathcal{ST}(G)$, meaning $\mathrm{w}(i,j) \ge \mathrm{r}(i,j)$ for every edge. For any w, there is a smallest constant $c > 0$ for which this property holds for $\tilde{\mathrm{w}} = c\mathrm{w}$; the below results yield the tightest risk bounds when applied to $\tilde{\mathrm{w}}$ scaled in this way. Whenever $\mathrm{w} \ge \mathrm{r}$ for some $\mathrm{r} \in \mathcal{ST}(G)$, (10) implies that the total weight of all edges satisfies

$$\mathrm{w}(E) \ge \mathrm{r}(E) = n - 1, \tag{11}$$

[4]We assume that G remains connected by the edges with strictly positive weight under w.

since every spanning tree has $n-1$ edges. The ratio $\mathrm{w}(E)/(n-1)$ provides a measure of the distance of w to $\mathcal{ST}(G)$. Furthermore, if E' is any subset of edges whose removal disconnects G into $k+1$ connected components, then every spanning tree contains at least k edges of E', so $\mathrm{w}(E') \geq \mathrm{r}(E') \geq k$. In particular,

$$\mu \in \mathbb{R}^n \text{ has } k+1 \text{ distinct values} \Rightarrow \|D\mu\|_{\mathrm{w}} \geq k. \tag{12}$$

We establish a sparsity oracle inequality when λ is set to a universal level:

Theorem 2.14. *Let* $\mathrm{w} : E \to \mathbb{R}_+$ *be such that* $\mathrm{w} \geq \mathrm{r}$ *for some* $\mathrm{r} \in \mathcal{ST}(G)$*, and let* $\delta \leq \sigma/\sqrt{n}$ *and* $\tau \leq \sigma^2$*. For any* $\eta > 0$*, there exist constants* $C_\eta, C'_\eta > 0$ *depending only on* η *such that if* $\lambda \geq C_\eta \sigma^2 \log \mathrm{w}(E)$ *and* $\hat{\mu}$ *is any* $(\tau, \delta\mathbb{Z})$*-local-minimizer of (W), then*

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \leq \inf_{\mu} (1+\eta)\|\mu - \mu_0\|^2 + C'_\eta \lambda \max(\|D\mu\|_{\mathrm{w}}, 1). \tag{13}$$

Evaluating at $\mu = \mu_0$ yields a risk bound of order $\sigma^2 s \log \mathrm{w}(E)$ when $\|D\mu_0\|_{\mathrm{w}} = s$. If s is known, then we may obtain a tighter guarantee by setting the penalty to a lower level:

Theorem 2.15. *Let* $\mathrm{w} : E \to \mathbb{R}_+$ *be such that* $\mathrm{w} \geq \mathrm{r}$ *for some* $\mathrm{r} \in \mathcal{ST}(G)$*, and let* $\delta \leq \sigma/\sqrt{n}$ *and* $\tau \leq \sigma^2$*. There exist universal constants* $C, C' > 0$ *such that for any* $s \in [1, \mathrm{w}(E)]$*, if* $\lambda \geq C\sigma^2(1 + \log \frac{\mathrm{w}(E)}{s})$ *and* $\hat{\mu}$ *is any* $(\tau, \delta\mathbb{Z})$*-local-minimizer of (W), then*

$$\sup_{\mu_0 : \|D\mu_0\|_{\mathrm{w}} \leq s} \mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \leq C' \lambda s. \tag{14}$$

Conversely, there exists a universal constant $c > 0$ *such that for any* $s \in [\frac{4\mathrm{w}(E)}{n}, \mathrm{w}(E)]$*,*

$$\inf_{\hat{\mu}} \sup_{\mu_0 : \|D\mu_0\|_{\mathrm{w}} \leq s} \mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \geq c\sigma^2 s \frac{n}{\mathrm{w}(E)} \left(1 + \log \frac{\mathrm{w}(E)}{s}\right),$$

where the infimum is taken over all possible estimators $\hat{\mu} := \hat{\mu}(Y)$*.*

The restriction to $s \gtrsim \mathrm{w}(E)/n$ in the lower bound is necessary for generality of the result to all graphs G, for the same reason as discussed in Remark 2.7.

The minimax upper and lower bounds above differ by the factor $n/\mathrm{w}(E)$. Recall from (11) that $\mathrm{w}(E) \geq n-1$, with $\mathrm{w}(E) = n-1$ precisely when $\mathrm{w} \in \mathcal{ST}(G)$. Hence the above immediately implies the following corollary:

Corollary 2.16. *If* $\mathrm{w} = \mathrm{r}$ *where* $\mathrm{r}(i,j)$ *is the effective resistance of edge* $\{i,j\}$*, or more generally where* r *is any edge-weighting in the spanning-tree polytope* $\mathcal{ST}(G)$*, then for any* $s \in [4, n-1]$*,*

$$\inf_{\hat{\mu}} \sup_{\mu_0 : \|D\mu_0\|_{\mathrm{w}} \leq s} \mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \asymp \sigma^2 s \left(1 + \log \frac{n}{s}\right). \tag{15}$$

Remark 2.17. While there is a strong analogy between (15) and the sparse normal-means result (5), it is perhaps less evident that the rate in (15) is correct.

Consider $\sigma^2 = 1$. In the sparse normal-means problem, if $\|\mu\|_0 \leq k$, then there are $\binom{n}{k} \leq e^{k(1+\log \frac{n}{k})}$ different possible sets of support for μ, i.e. at most $e^{k(1+\log \frac{n}{k})}$ linear sub-models $\mathcal{M}$ to which μ belongs. Penalizing each non-zero entry of μ with penalty $\lambda \asymp 1 + \log \frac{n}{k}$ exactly balances this complexity in the number of sub-models, and hence a risk bound of $Ck \log(1 + \frac{n}{k})$ over the class $\{\mu_0 : \|\mu_0\|_0 \leq k\}$ follows from, e.g., the general analysis of model-complexity penalties in [BBM99].

In contrast, consider in our setting the complete graph G on n vertices, with w given by effective resistance weighting and $s \asymp n/\log n$. Consider all partitions of the vertices of G into $l := n/\log n$ sets $S_1, \ldots, S_l$ of two vertices each, plus one set S_{l+1} of the remaining $n - 2l$ vertices. Each such partition defines a linear sub-model $\mathcal{M}$ such that every $\mu \in \mathcal{M}$ takes a constant value over each vertex set $S_1, \ldots, S_{l+1}$. One way to construct such a partition assigns $1 \in S_1, 2 \in S_2, \ldots, l \in S_l$, and then picks one additional vertex to add to each of $S_1, \ldots, S_l$. The number of such choices, and

hence the number of such sub-models, is greater than $(n-2l)^l = e^{l\log n - o(l\log n)} = e^{n-o(n)}$. On the other hand, for any μ belonging to such a sub-model, the number of edges $\{i,j\}$ for which $\mu_i \neq \mu_j$ is of order $nl = n^2/\log n$. Each edge has effective resistance $\mathrm{w}(i,j) = 2/n$ in the complete graph, so $\|D\mu_0\|_{\mathrm{w}} \leq s$. Setting $\lambda \asymp 1 + \log\frac{n}{s} \asymp \log\log n$ penalizes each such sub-model by a penalty of $\lambda\|D\mu_0\|_{\mathrm{w}} \lesssim (n\log\log n)/\log n$, which is not enough to balance the number of sub-models, $e^{n-o(n)}$. Naively, it seems like the penalty $\lambda \asymp 1 + \log\frac{n}{s}$ is too small to yield the risk bound (14).

We circumvent this problem in Lemma A.2 by noting that even though the number of sub-models is large, the vertex partitions corresponding to many sub-models are refined by a single common refinement. We obtain the upper bound (14) by taking a union bound over only the set of common refinements, whose cardinality is of a smaller size.

Remark 2.18. One may compare (14) with a guarantee achieved by the wavelet spanning tree method of [SSR13]: In this method, for some spanning tree T, an orthonormal basis of wavelet functions is constructed over T such that a signal μ_0 cutting only a few edges $\{i,j\} \in T$ has a sparse representation in this basis. Theorem 56 of [Sha13] implies that the thresholding estimator $\hat{\mu}$ in this basis satisfies

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \lesssim \sigma^2(\log n)(\log d_{\max}(T))|\{\{i,j\} \in T : \mu_{0,i} \neq \mu_{0,j}\}|,$$

where $d_{\max}(T)$ is the maximum vertex degree of T. If T is chosen at random from the spanning tree distribution corresponding to any weighting $\mathrm{w} \in \mathcal{ST}(G)$, then bounding $d_{\max}(T) \leq d_{\max}(G)$ and then averaging over T yields

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \lesssim \sigma^2(\log n)(\log d_{\max}(G))\|D\mu_0\|_{\mathrm{w}},$$

which is similar to (14). Whereas this defines a randomized algorithm and the above risk is averaged also over the algorithm execution, minimizing (W) for $\mathrm{w} \in \mathcal{ST}(G)$ directly penalizes the number of edges cut by $\hat{\mu}$ in the average spanning tree, and thus may be interpreted loosely as a derandomization of this wavelet approach.

Finally, let us establish a result of relevance to testing the null hypothesis (7):

Theorem 2.19. *Let* $\mathrm{w} : E \to \mathbb{R}_+$ *be such that* $\mathrm{w} \geq \mathrm{r}$ *for some* $\mathrm{r} \in \mathcal{ST}(G)$*, and let* $\delta \leq \sigma/\sqrt{n}$ *and* $\tau \leq \sigma^2$*. There exist universal constants* $C, C' > 0$ *such that if* μ_0 *is constant over* G*,* $\lambda \geq C\sigma^2 \log \mathrm{w}(E)$*, and* $\hat{\mu}$ *is any* $(\tau, \delta\mathbb{Z})$*-local-minimizer of (W), then*

$$\mathbb{P}[\hat{\mu} \text{ is constant over } G] \geq 1 - C'n^{-3}.$$

Thus, for any weighting $\mathrm{w} : E \to \mathbb{R}_+$, we may test the null hypothesis (7) by setting $\lambda \asymp \sigma^2 \log \mathrm{w}(E)$ in the weighted objective function (W) and rejecting H_0 if $\hat{\mu}$ minimizing (W) is not constant. Denoting by $P^\perp$ the projection orthogonal to the all-1's vector, the risk bound (14) (or more precisely, the probability guarantee of Lemma A.3(b)) implies that this test can distinguish a non-constant alternative μ_0 with high probability as long as $\|P^\perp\mu_0\|^2 \gtrsim \sigma^2\|D\mu_0\|_{\mathrm{w}} \log \mathrm{w}(E)$. When $\mathrm{w} : E \to \mathbb{R}_+$ is the effective resistance weighting, this recovers a similar detection threshold as established for the tests in [SSR13, SKS13].

We prove Theorems 2.14, 2.15, and 2.19 in Appendix A. In the case of uniform edge weights $\mathrm{w} \equiv 1$, it is clear that $\mathrm{w}(E) = |E|$ and $\mathrm{w} \geq \mathrm{r}$ for all $\mathrm{r} \in \mathcal{ST}(G)$. Then Theorems 2.4, 2.5, 2.6, and 2.8 for the unweighted objective (L0) follow directly by specializing these results. If there exists $\mathrm{r} \in \mathcal{ST}(G)$ such that $\mathrm{r}(i,j) < 1$ for every edge, then the results of Section 2.2 are trivially strengthened by rescaling λ by $\max_{\{i,j\}\in E} \mathrm{r}(i,j)$. For example, if G is the complete graph, then every edge has effective resistance $\mathrm{r}(i,j) = 2/n$, and Theorems 2.14 and 2.15 imply that λ may in fact be set to $Cn^{-1}\log|E|$ and $Cn^{-1}\log|E|/s$ in Theorems 2.4 and 2.5 respectively, as claimed in Example 2.11.

3. Simulations

We study empirically the squared-errors of the approximate minimizers of (L0) and (W) as returned by Algorithm 1, as well as the exact minimizer of (TV). We will denote these estimates by $\hat{\mu}^{\mathrm{L0}}$, $\hat{\mu}^{\mathrm{W}}$, and $\hat{\mu}^{\mathrm{TV}}$. We consider piecewise-constant signals over the linear chain, 2-D lattice, and several real-world networks, corrupted by artificial Gaussian noise for various noise levels σ. We report in each setting the standardized mean-squared-error

$$\frac{1}{n\sigma^2}\|\hat{\mu} - \mu_0\|^2. \tag{st.MSE}$$

Due to this normalization by σ^2, one may equivalently interpret these results as for a fixed noise level σ under various rescalings of the true signal μ_0.

3.1. **Parameter tuning.** For Algorithm 1, we fix throughout $\delta = 0.01$ and $\tau = 0$. This value of δ is not necessarily that prescribed by the theory of the preceding section, but represents a compromise to yield faster runtime. We replace the estimated value for each constant piece of $\hat{\mu}^{\mathrm{L0}}$ and $\hat{\mu}^{\mathrm{W}}$ by the sample mean of Y over this piece (which differs by at most δ), to mitigate some of the effect of this discretization.

We select the penalty λ by minimizing an empirical estimate of $\mathbb{E}[\|\hat{\mu} - \mu_0\|^2]$. Typically, cross-validation is used to obtain such an estimate. However, we observe that naive cross-validation does not necessarily work well for all graphs and signals. (Consider, for example, a case where the primary contribution to error comes from vertices i near the boundaries of the constant pieces of μ_0, and estimation of these values $\mu_{0,i}$ is more difficult when Y_i is removed.) We instead use the following procedure based on [TT15, Har16]:

(1) Compute an estimate $\hat{\sigma}$ for σ. Set $\alpha = 0.04$.
(2) For repetitions $b = 1, \dots, B$:
    (a) Generate $z = (z_1, \dots, z_n) \sim \mathrm{Normal}(0, \alpha\hat{\sigma}^2\,\mathrm{Id})$, and set $Y^* = Y + z$ and $Y^{**} = Y - z/\alpha$.
    (b) For each λ, compute $\hat{\mu}$ based on data Y^*, and compute $\mathrm{err}^{(b)}(\lambda) = \|\hat{\mu} - Y^{**}\|^2$.
(3) Choose λ that minimizes the average error $\overline{\mathrm{err}}(\lambda) = \frac{1}{B}\sum_{b=1}^{B} \mathrm{err}^{(b)}(\lambda)$.

This is motivated by the insight that if $\hat{\sigma} = \sigma$, then Y^* and Y^{**} are independent, so $\mathbb{E}[\mathrm{err}^{(b)}(\lambda)] = n\sigma^2(1 + \alpha^{-1}) + \mathbb{E}[\|\hat{\mu} - \mu_0\|^2]$. Hence $\overline{\mathrm{err}}(\lambda)$ estimates a constant plus the risk of $\hat{\mu}$ applied to data at the slightly elevated noise level $\sigma(1 + \sqrt{\alpha}) = 1.2\sigma$. Due to this elevation in noise level, this procedure has a slight tendency to oversmooth.

Note that for each edge $\{i, j\}$ where $\mu_{0,i} \neq \mu_{0,j}$, $Y_i - Y_j \sim \mathrm{Normal}(0, 2\sigma^2)$. Hence σ may be estimated from the edge differences $(Y_i - Y_j)_{\{i,j\}\in E}$ by identifying a normal mixture component corresponding to this subset of values; we used the method of [Efr04] as implemented in the locfdr R package and found results to be very accurate in all tested settings. Increasing B reduces the variability of the selection procedure. For the smaller graphs (linear chain, Oldenburg, Gnutella P2P) we set $B = 20$, and for the larger graphs (2-D cow, San Francisco, Enron email) we set $B = 5$.

We will report both the st.MSE achieved using this method, as well as the best-attained st.MSE corresponding to retrospective optimal tuning of λ. For (TV), λ may alternatively be selected by minimizing Stein's unbiased risk estimate (SURE) using the simple degrees-of-freedom formula derived in [TT11]. We tested this approach and found results to be extremely close to those obtained using the above procedure, so we will omit the SURE results for consistency and brevity.

3.2. **Empirical runtime.** For Algorithm 1, we used the max-flow code of Boykov and Kolmogorov, based on [BK04], to compute the minimum s-t cut. On our personal computer, average runtime for computing $\hat{\mu}^{\mathrm{L0}}$ or $\hat{\mu}^{\mathrm{W}}$ (for a single λ) was approximately 0.13 seconds for the linear chain graph with $n = 1000$ vertices, 45 seconds for the cow image on a 2-D lattice graph of $n \approx 90000$ vertices, and 0.7 seconds, 40 seconds, 4 seconds, and 4 minutes for the Oldenburg, San Francisco, Gnutella P2P, and Enron email networks, respectively. (This runtime for $\hat{\mu}^{\mathrm{W}}$ does not include computation

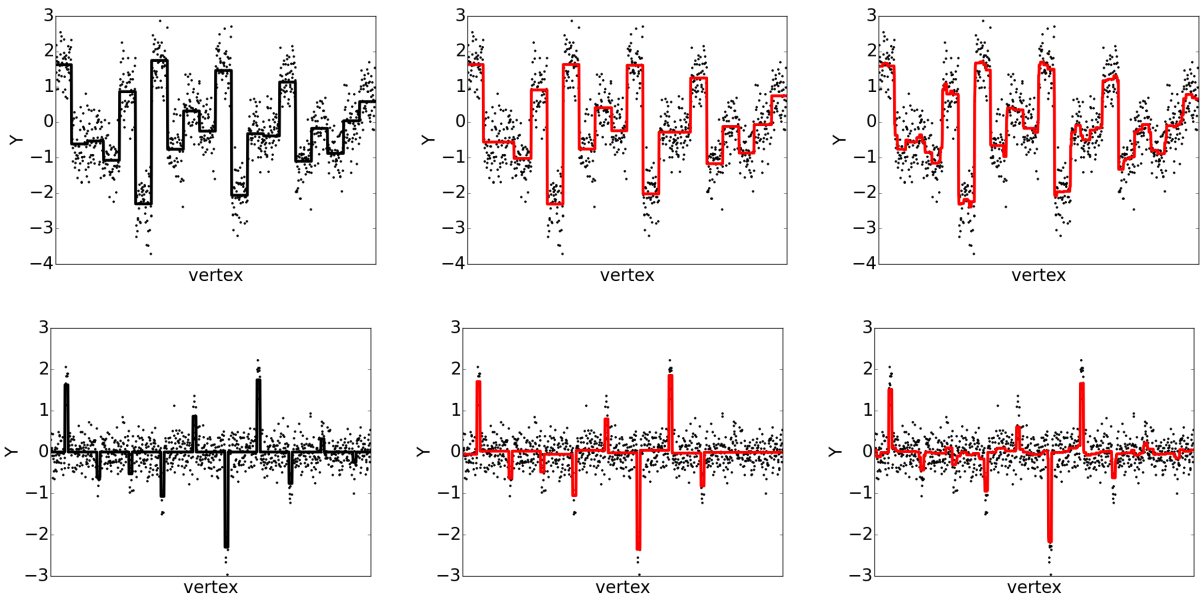

FIGURE 1. Signals on a linear chain graph. Top: Equally-spaced breaks, $\sigma = 0.5$. Bottom: Unequally-spaced breaks, $\sigma = 0.3$. The true signal μ_0 is displayed on the left, $\hat{\mu}^{\mathrm{L0}}$ in the middle, and $\hat{\mu}^{\mathrm{TV}}$ on the right (both with data-tuned λ).

of the effective edge resistances, which in our experience can be much slower. We computed these resistances using the Lean Algebraic MultiGrid Matlab software based on [LB12].)

For computing $\hat{\mu}^{\mathrm{TV}}$, we used the pygfl library of Tansey and Scott, based on [TS15]. Average runtime for computing $\hat{\mu}^{\mathrm{TV}}$ without algorithm tuning was approximately 0.04 seconds for the linear chain graph, 25 seconds for the cow image, and 0.9 seconds, 90 seconds, 0.5 seconds, and 7 seconds for the Oldenburg, San Francisco, Gnutella P2P, and Enron email networks. (Alternative algorithms can yield faster runtime for 2-D images.)

Parameter tuning using the approach of Section 3.1 represents a substantial increase in computational cost, although the additional computation is easily parallelized.

3.3. **Linear chain graph.** Two signals on a linear chain graph with $n = 1000$ vertices are depicted in Figure 1. The first signal has 19 equally-spaced break points, while the second has 20 break points at unequal spacing. (In the equally-spaced case, $\Delta(\mu_0) = n/(s+1)$ so that the lower bound of Theorem 2.10 does not necessarily imply suboptimality of the total-variation approach even in a minimax sense.) We studied recovery of these signals for noise levels $\sigma = 0.1$ to $\sigma = 1$. Figure 1 displays one instance of simulated noise and the resulting estimates $\hat{\mu}^{\mathrm{L0}}$ and $\hat{\mu}^{\mathrm{TV}}$. In both examples, for data-tuned λ, we observe the tendency of $\hat{\mu}^{\mathrm{L0}}$ to over-smooth (missing two and four changepoints respectively) and $\hat{\mu}^{\mathrm{TV}}$ to undersmooth.

Figure 2 plots st.MSE comparisons for $\hat{\mu}^{\mathrm{L0}}$ and $\hat{\mu}^{\mathrm{TV}}$. In this linear chain example, it is interesting to also compare $\hat{\mu}^{\mathrm{L0}}$ with the exact minimizer of (L0), which we computed using the PELT method in the changepoint R package [KFE12]. We in fact observe that Algorithm 1 achieves risk comparable to that of the exact minimizer in all tested settings. $\hat{\mu}^{\mathrm{L0}}$ achieves significantly smaller risk than $\hat{\mu}^{\mathrm{TV}}$ at higher signal-to-noise regimes, for example those displayed in Figure 1, while $\hat{\mu}^{\mathrm{TV}}$ becomes competitive or better in lower signal-to-noise regimes (corresponding to lower values of $\|D\mu_0\|_1/\sigma$ for the true signal). The improvement of $\hat{\mu}^{\mathrm{L0}}$ over $\hat{\mu}^{\mathrm{TV}}$ is greater in the unequally-spaced setting, which is in qualitative agreement with the theoretical comparisons of Section 2.3.

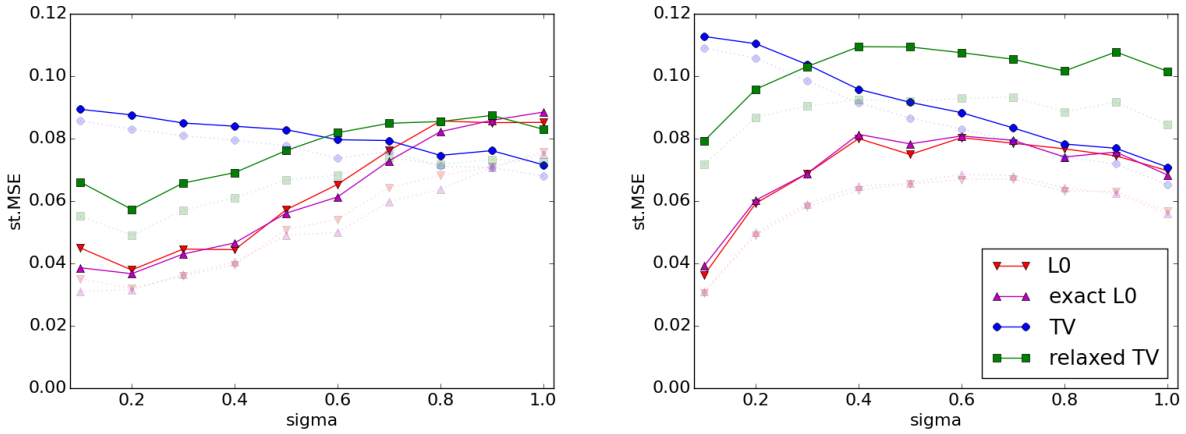

FIGURE 2. Comparisons of st.MSE in (left) the equally-spaced and (right) unequally-spaced examples of Figure 1, for $\hat{\mu}^{\mathrm{L0}}$, $\hat{\mu}^{\mathrm{TV}}$, the exact minimizer of (L0), and the relaxed/debiased version of $\hat{\mu}^{\mathrm{TV}}$. Solid lines correspond to data-tuned λ, and dashed transparent lines to best-achieved error. All errors are averaged over 100 replicates of the simulated noise.

One may ask, at the higher signal-to-noise regimes, how much of the sub-optimality of $\hat{\mu}^{\mathrm{TV}}$ is due to the estimator bias incurred by shrinkage. To address this question, we computed also the "relaxed" TV estimate obtained by replacing each constant interval of $\hat{\mu}^{\mathrm{TV}}$ with the mean of Y over this interval. The st.MSE of this estimate is also displayed in Figure 2—the sub-optimality at high signal-to-noise is partially (but not completely) ameliorated, but the improvement at lower signal-to-noise is largely lost.

3.4. **2-D lattice graph.** Figure 3 displays a cartoon gray-scale image of a cow, represented by its pixel values on a 2-D lattice graph of size 320×283. Pure white corresponds to $\mu_0 = 1$, and pure black to $\mu_0 = 0$. The figure also displays $\hat{\mu}^{\mathrm{L0}}$ and $\hat{\mu}^{\mathrm{TV}}$ when the image is contaminated by noise at level $\sigma = 0.3$. We again observe that $\hat{\mu}^{\mathrm{L0}}$ oversmooths, missing details in the cow's feet, right horn, and the shadows of the image. In contrast, $\hat{\mu}^{\mathrm{TV}}$ undersmooths and returns a blotchy cow.

Table 1 reports st.MSE comparisons for $\sigma = 0.1$ to $\sigma = 0.5$. At the level $\sigma = 0.3$ displayed in Figure 3, the st.MSE of $\hat{\mu}^{\mathrm{L0}}$ is slightly greater than that of $\hat{\mu}^{\mathrm{TV}}$. At higher signal-to-noise levels, $\hat{\mu}^{\mathrm{L0}}$ is better, while it is worse at lower signal-to-noise.

3.5. **Road and digital networks.** We tested signal recovery over four real networks: the Oldenburg and San Francisco road networks from `www.cs.utah.edu/~lifeifei/SpatialDataset.htm`, and the Gnutella08 peer-to-peer network and Enron email network from `snap.stanford.edu/data`.

sigma	0.1	0.2	0.3	0.4	0.5
$\hat{\mu}^{\mathrm{L0}}$ st.MSE	0.041 (0.039)	0.067 (0.067)	0.076 (0.071)	0.081 (0.081)	0.082 (0.082)
$\hat{\mu}^{\mathrm{TV}}$ st.MSE	0.083 (0.083)	0.075 (0.075)	0.065 (0.065)	0.061 (0.059)	0.054 (0.053)

TABLE 1. Comparison of st.MSE for the cow image of Figure 3. Non-parenthesized values correspond to data-tuned λ, and parenthesized values to best-attained error.

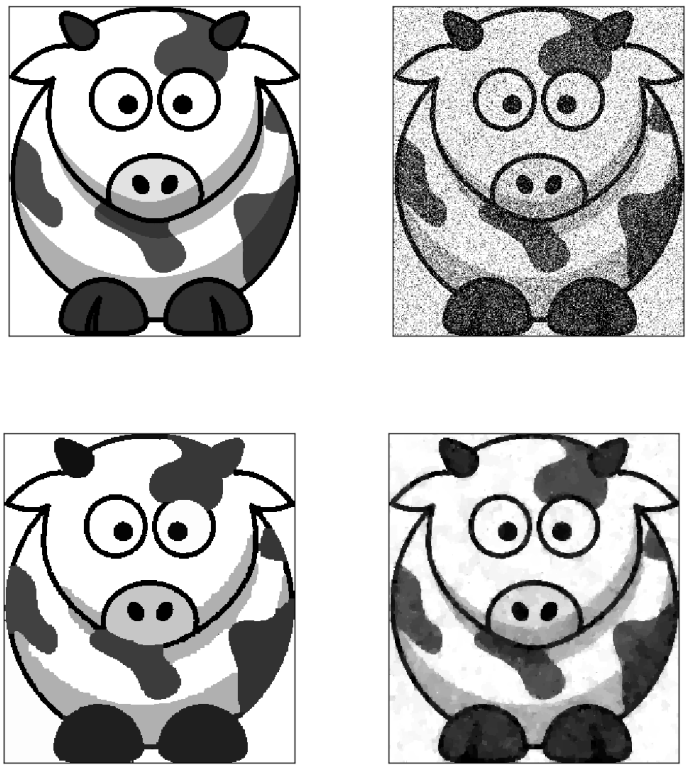

FIGURE 3. (top left) Original image, with pixel values normalized to $[0, 1]$. (top right) Noisy image, $\sigma = 0.3$. (bottom left) $\hat{\mu}^{\mathrm{L0}}$ and (bottom right) $\hat{\mu}^{\mathrm{TV}}$, both with data-tuned λ.

Duplicate edges were removed, and only the largest connected component of each network was retained.

For each network, we simulated an epidemic according to a simple graph-based discrete-time SI model [MN00], by randomly selecting a source vertex to infect at time $t = 0$ and, for each of T timesteps, allowing each infected vertex to independently infect each non-infected neighbor with probability 0.5. Over the digital networks, this may represent the spread of an internet/email virus. Over the road networks, this may represent the spread of zombie infections in a zombie apocalypse. We associated to each infected vertex the value $\mu_0 = 1$, and to each non-infected vertex

Network	verts.	edges	res. var.	inf. T_1	cut T_1	inf. T_2	cut T_2	inf. T_3	cut T_3
Oldenburg	6105	7029	0.118	57	16	515	75	2108	164
San Francisco	174956	221802	0.203	8574	306	27724	562	70925	774
Gnutella P2P	6299	20776	0.826	19	91	67	477	271	3894
Enron email	33696	180811	1.297	179	7319	2564	74117	16868	29253

TABLE 2. For each network: Number of total vertices, total edges, variability of effective edge resistances (measured by standard deviation / mean), and numbers of infected vertices and cut edges corresponding to the signal at three observation times.

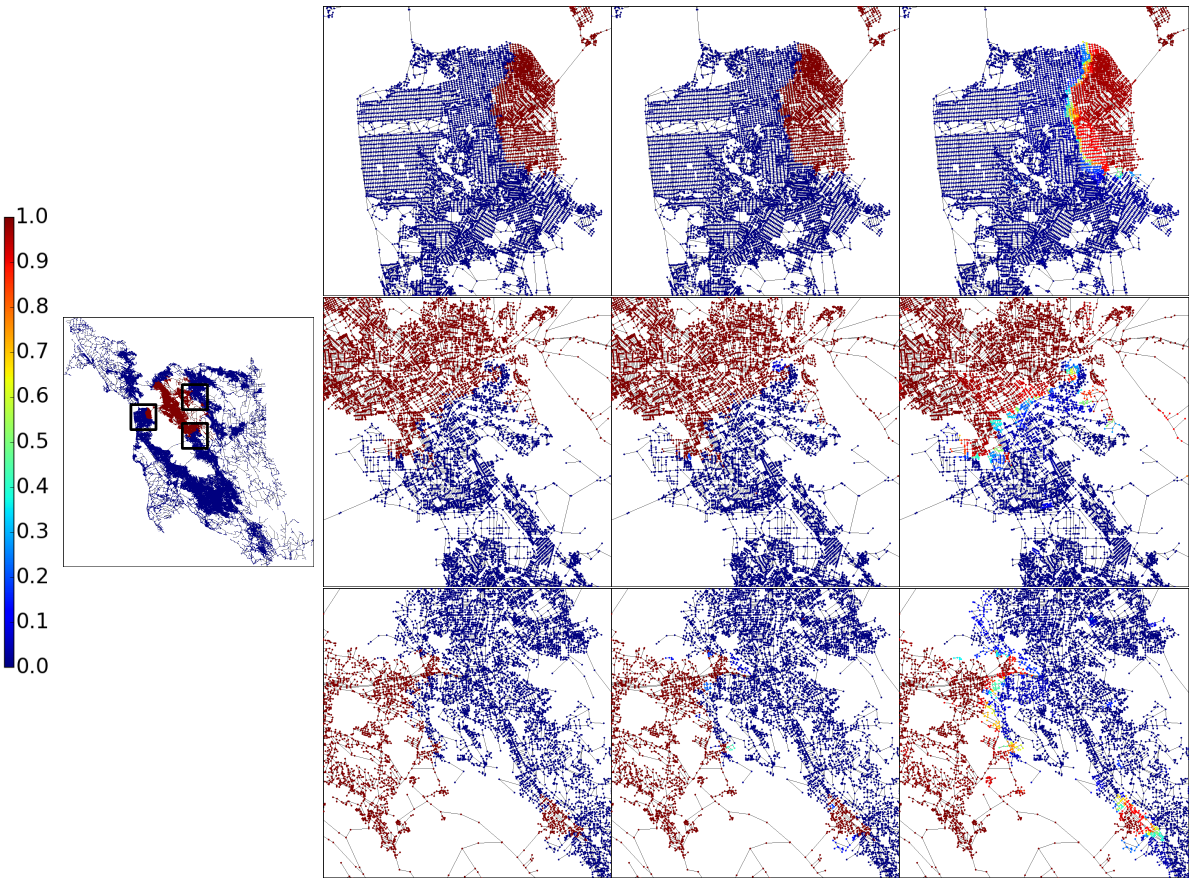

FIGURE 4. A simulated epidemic signal over the San Francisco road network is displayed on the far left. Three columns display (left) the true signal, (middle) $\hat{\mu}^{\mathrm{W}}$, and (right) $\hat{\mu}^{\mathrm{TV}}$ for the three boxed areas of the map. Both $\hat{\mu}^{\mathrm{W}}$ and $\hat{\mu}^{\mathrm{TV}}$ are estimated on noisy data with $\sigma = 0.3$, using data-tuned λ, and $\hat{\mu}^{\mathrm{W}}$ uses effective-resistance edge weighting.

the value $\mu_0 = 0$.[5] For each network, we considered three signals corresponding to observations of the epidemic at three different times T. Various properties of these networks and signals are summarized in Table 2.

For noise level $\sigma = 0.3$, the simulated signal, $\hat{\mu}^{\mathrm{TV}}$, and $\hat{\mu}^{\mathrm{W}}$ computed with effective-resistance edge weights are depicted in Figure 4 for the San Francisco road network at observation time T_2. The most difficult regions to estimate in this example are the boundaries of the zombie infection; we have zoomed in on three regions of the map where $\hat{\mu}^{\mathrm{TV}}$ does not yield accurate estimation at these boundaries, but $\hat{\mu}^{\mathrm{W}}$ is mostly correct. At this noise level, the st.MSE of $\hat{\mu}^{\mathrm{TV}}$ exceeds that of $\hat{\mu}^{\mathrm{W}}$ by a factor of about 2.

Figure 5 displays st.MSE comparisons for $\hat{\mu}^{\mathrm{L0}}$, $\hat{\mu}^{\mathrm{TV}}$, and $\hat{\mu}^{\mathrm{W}}$ with effective-resistance weighting at noise levels $\sigma = 0.1$ to $\sigma = 0.5$ in each example. We observe that $\hat{\mu}^{\mathrm{W}}$ is not substantially worse than $\hat{\mu}^{\mathrm{L0}}$ in any tested setting, and that in the Gnutella and Enron digital networks where there is large variation in effective edge resistances, $\hat{\mu}^{\mathrm{W}}$ is sometimes substantially better. At the noise

[5]Actually, we assigned values 0.005 and 1.005, to avoid possible artificial benefits given to Algorithm 1 under the discretization $\delta = 0.01$. In this simple example, it is clear that methods with a priori knowledge that there are only two distinct signal values will do better; we do not use this information in the methods that we test.

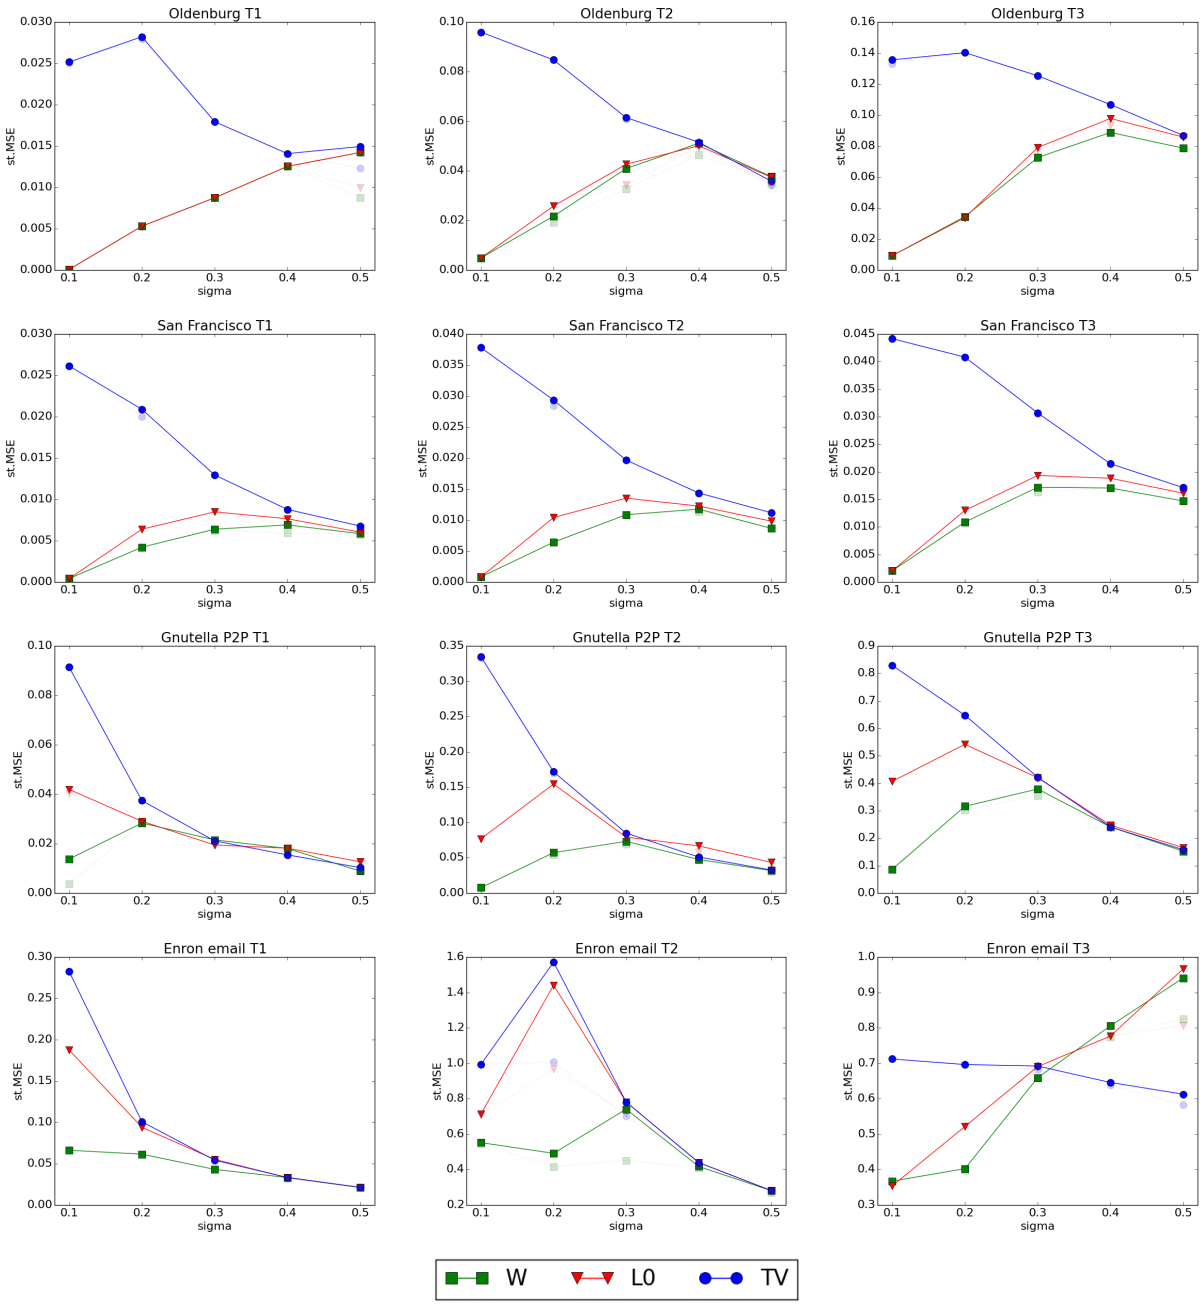

FIGURE 5. Comparisons of st.MSE for recovery of epidemic signals on four networks, for $\hat{\mu}^{\mathrm{L0}}$, $\hat{\mu}^{\mathrm{TV}}$, and $\hat{\mu}^{\mathrm{W}}$ with effective-resistance edge weights. Solid lines correspond to data-tuned λ, and dashed transparent lines to best-achieved error.

levels that we tested, these methods are in most cases (with the exception of Enron T_3) also not substantially worse than $\hat{\mu}^{\mathrm{TV}}$, and can be substantially better in the lower noise settings.

4. Conclusion

We have studied the estimation of piecewise-constant signals over arbitrary graphs, using the method of l_0-denoising. We have provided theoretical evidence that the α-expansion procedure of [BVZ01] may be applied to yield a good approximate solution to the l_0 problem for any graph, in polynomial time. Our problem may be reformulated as a sparse regression problem with a linear constraint on the regression coefficients and a particular graph-based design; we believe it is an interesting question whether computational ideas from discrete Markov random fields may be used to achieve better prediction error in sparse regression for other more general families of designs.

For application to networks with regions of varying connectivity, we have recommended the minimization of an edge-weighted l_0 objective. In addition to achieving better performance in some tested examples, a spanning-tree-based choice of edge weights leads to a simple proof of theoretical guarantees that does not require detailed understanding of the graph structure.

In many examples, we have found that approximate minimization of the l_0 objective can yield better signal recovery than exact minimization of the l_1 relaxation, with reasonable runtime. We recommend that l_0-based methods be used in high signal-to-noise settings, when rough recovery of the constant signal regions is possible. In noisier settings and when edge-sparsity of the recovered signal is not of primary interest, l_1-based methods provide a possibly better alternative.

Appendix A. Analysis of weighted l_0-denoising

In this appendix, we prove Theorems 2.14, 2.15, and 2.19. By rescaling, we will assume without loss of generality that $\sigma^2 = 1$.

Terminology. A **multi-cut** $\mathbf{S} := \{S_1, \dots, S_k\}$ of G is a partition of its vertices V into non-empty disjoint vertex subsets $S_1, \dots, S_k$, which are called the **elements** of the multi-cut. In this appendix, we do not require each element S_α of $\mathbf{S}$ to be connected in G. The multi-cut $\mathbf{S}$ is **trivial** if $k = 1$, i.e. its only element is the entire vertex set V. For two multi-cuts $\mathbf{S}$ and $\mathbf{S}'$, $\mathbf{S}'$ is a **refinement** of $\mathbf{S}$ if this holds in the usual sense of partitions: For every element S_α of $\mathbf{S}'$, there exists an element of $\mathbf{S}$ that contains S_α as a subset. For any two multi-cuts $\mathbf{S}$ and $\mathbf{S}'$, their least common refinement is denoted by $\mathbf{S} \vee \mathbf{S}'$.

For every $\mu \in \mathbb{R}^n$, the multi-cut **induced by** μ is such that any two vertices $i, j \in V$ are in the same element of the cut if and only if $\mu_i = \mu_j$. (If μ is constant on G, then it induces the trivial multi-cut.) A multi-cut **cuts** an edge $\{i, j\}$ if i and j belong to two different elements. The set of all cut edges of $\mathbf{S}$ is denoted

$$\partial \mathbf{S} := \{\{i, j\} \in E : \mathbf{S} \text{ cuts } \{i, j\}\}.$$

($\partial \mathbf{S}$ is empty if $\mathbf{S}$ is trivial.) Recall $\mathrm{w}(E') := \sum_{\{i,j\} \in E'} \mathrm{w}(i, j)$ for any $E' \subseteq E$. It is then evident that if μ induces $\mathbf{S}$, then

$$\|D\mu\|_{\mathrm{w}} = \mathrm{w}(\partial \mathbf{S}).$$

Associated to each multi-cut $\mathbf{S} = \{S_1, \dots, S_k\}$ is a k-dimensional subspace $K \subseteq \mathbb{R}^n$, such that $\mu \in K$ if and only if μ takes a constant value on each element S_α. (I.e., $\mu \in K$ if and only if $\mathbf{S}$ is equal to, or is a refinement of, the multi-cut induced by μ.) We denote the orthogonal projection onto K by $P^{\mathbf{S}} : \mathbb{R}^n \to K$.

We first recall a deterministic sub-optimality bound for any $(\tau, \delta\mathbb{Z})$-local-minimizer of (W). Such a bound was stated in [BVZ01] as a factor-2 approximation, $F_{\mathrm{w}}(\hat{\mu}) \leq 2F_{\mathrm{w}}(\mu^*)$, where $\hat{\mu}$ is any $(0, \delta\mathbb{Z})$-local-minimizer of (W) and μ^* is the exact minimizer in $(\delta\mathbb{Z})^n$. Examination of the proof reveals that this factor of 2 applies only to the penalty term, the resulting bound holds not only for μ^* but also for any $\mu \in (\delta\mathbb{Z})^n$, and it is easily extended to any $\tau \geq 0$—we restate the argument for completeness.

Lemma A.1. *Let $\delta > 0$ and $\tau \geq 0$, and let $\hat{\mu}$ be any $(\tau, \delta\mathbb{Z})$-local-minimizer of $F_{\mathrm{w}}(\mu)$ in (W). Then for any vector $\mu^\delta \in (\delta\mathbb{Z})^n$ whose entries take k distinct values,*

$$F_{\mathrm{w}}(\hat{\mu}) \leq \frac{1}{2}\|Y - \mu^\delta\|^2 + 2\lambda\|D\mu^\delta\|_{\mathrm{w}} + k\tau.$$

Proof. Let $\mathbf{S} = \{S_1, \ldots, S_k\}$ be the multi-cut induced by μ^δ. Fix $\alpha \in \{1, \ldots, k\}$ and denote by $\mu^\delta_{S_\alpha}$ the (constant) value of μ^δ on S_α. Let $\tilde{\mu}$ be the $\delta\mathbb{Z}$-expansion of $\hat{\mu}$ defined entrywise by

$$\tilde{\mu}_i = \begin{cases} \hat{\mu}_i & i \notin S_\alpha \\ \mu^\delta_{S_\alpha} & i \in S_\alpha. \end{cases}$$

Since $\hat{\mu}$ is a $(\tau, \delta\mathbb{Z})$-local-minimizer of (W), by definition

$$\begin{aligned} F_{\mathrm{w}}(\hat{\mu}) &= \frac{1}{2}\sum_{i=1}^n (Y_i - \hat{\mu}_i)^2 + \lambda \sum_{\{i,j\}\in E} \mathrm{w}(i,j)\mathbb{1}\{\hat{\mu}_i \neq \hat{\mu}_j\} \\ &\leq \frac{1}{2}\sum_{i=1}^n (Y_i - \tilde{\mu}_i)^2 + \lambda \sum_{\{i,j\}\in E} \mathrm{w}(i,j)\mathbb{1}\{\tilde{\mu}_i \neq \tilde{\mu}_j\} + \tau = F_{\mathrm{w}}(\tilde{\mu}) + \tau. \end{aligned}$$

As $\hat{\mu}_i = \tilde{\mu}_i$ for $i \notin S_\alpha$, we may cancel terms from both sides to yield

$$\frac{1}{2}\sum_{i\in S_\alpha}(Y_i - \hat{\mu}_i)^2 + \lambda \sum_{\substack{\{i,j\}\in E \\ |\{i,j\}\cap S_\alpha|\geq 1}} \mathrm{w}(i,j)\mathbb{1}\{\hat{\mu}_i \neq \hat{\mu}_j\} \leq \frac{1}{2}\sum_{i\in S_\alpha}(Y_i - \mu^\delta_{S_\alpha})^2 + \lambda \sum_{\substack{\{i,j\}\in E \\ |\{i,j\}\cap S_\alpha|\geq 1}} \mathrm{w}(i,j)\mathbb{1}\{\tilde{\mu}_i \neq \tilde{\mu}_j\} + \tau.$$

For the right side above, note that $\mathbb{1}\{\tilde{\mu}_i \neq \tilde{\mu}_j\} = 0$ when $i, j \in S_\alpha$, and apply the trivial bound $\mathbb{1}\{\tilde{\mu}_i \neq \tilde{\mu}_j\} \leq 1$ when $|\{i,j\} \cap S_\alpha| = 1$, i.e. when $\{i,j\}$ is cut by $\mathbf{S}$. Then, summing the above over $\alpha = 1, \ldots, k$,

$$\begin{aligned} \frac{1}{2}\|Y - \hat{\mu}\|^2 + \lambda \sum_{\alpha=1}^k \sum_{\substack{\{i,j\}\in E \\ |\{i,j\}\cap S_\alpha|\geq 1}} \mathrm{w}(i,j)\mathbb{1}\{\hat{\mu}_i \neq \hat{\mu}_j\} &\leq \frac{1}{2}\|Y - \mu^\delta\|^2 + \lambda \sum_{\alpha=1}^k \sum_{\substack{\{i,j\}\in E \\ |\{i,j\}\cap S_\alpha|=1}} \mathrm{w}(i,j) + k\tau \\ &= \frac{1}{2}\|Y - \mu^\delta\|^2 + 2\lambda\|D\mu^\delta\|_{\mathrm{w}} + k\tau, \end{aligned}$$

where the second line follows from noting that each edge $\{i,j\}$ cut by $\mathbf{S}$ appears twice in the sum. The desired result then follows by noting that each edge $\{i,j\} \in E$ appears at least once in the double sum on the left side. □

Next, we control $\|P^{\mathbf{S}}\varepsilon\|^2$ using the weighted cut-size $\mathrm{w}(\partial\mathbf{S})$ uniformly over all multi-cuts $\mathbf{S}$ of G. As the number of distinct multi-cuts of G may be very large, we will first condition on a random spanning tree T of G, consider the common refinement of those multi-cuts that cut the same set of edges of T, and take a union bound over these refinements. The distribution of the random spanning tree T that we take is the one specified by the edge-weighting $\mathrm{r} \in \mathcal{ST}(G)$ such that $\mathrm{w} \geq \mathrm{r}$.

Lemma A.2. *Let $\mathrm{w} : E \to \mathbb{R}_+$ be such that $\mathrm{w} \geq \mathrm{r}$ for some $\mathrm{r} \in \mathcal{ST}(G)$, and let $\varepsilon \in \mathbb{R}^n$ have coordinates $\varepsilon_i \overset{iid}{\sim} \mathrm{Normal}(0,1)$. Then there exist universal constants $C, C' > 0$ such that with probability at least $1 - C'\mathrm{w}(E)^{-3}$, for every multi-cut $\mathbf{S}$ of G,*

$$\|P^{\mathbf{S}}\varepsilon\|^2 \leq C\max(\mathrm{w}(\partial\mathbf{S}), 1)\log\frac{e\cdot\mathrm{w}(E)}{\max(\mathrm{w}(\partial\mathbf{S}), 1)}.$$

Proof. If $\mathbf{S}$ is the trivial multi-cut consisting of the single element $\{1, \ldots, n\}$, then $\mathrm{w}(\partial\mathbf{S}) = 0$ and $\|P^{\mathbf{S}}\varepsilon\|^2 \sim \chi^2_1$, so $\|P^{\mathbf{S}}\varepsilon\|^2 \leq C\log(e\cdot\mathrm{w}(E))$ with probability at least $1 - C'\mathrm{w}(E)^{-3}$ by a standard Gaussian tail bound.

Consider, then, non-trivial multi-cuts $\mathbf{S}$. By (12), $\mathrm{w}(\partial\mathbf{S}) \geq 1$. Observe that for any $c > 0$, the function $x \mapsto x\log(ec/x)$ is increasing over $x \in (0, c)$. Then, as $\mathrm{w} \geq \mathrm{r}$,

$$C\mathrm{w}(\partial\mathbf{S}) \log \frac{e \cdot \mathrm{w}(E)}{\mathrm{w}(\partial\mathbf{S})} \geq C\mathrm{r}(\partial\mathbf{S}) \log \frac{e \cdot \mathrm{w}(E)}{\mathrm{r}(\partial\mathbf{S})},$$

and it suffices to establish the stronger bound given by the right side above.

For any $\eta > 0$, we may bound

$$\begin{aligned}
p_n &:= \mathbb{P}\left[\text{ there exists a non-trivial multi-cut } \mathbf{S} : \|P^{\mathbf{S}}\varepsilon\|^2 > C\mathrm{r}(\partial\mathbf{S}) \log \frac{e \cdot \mathrm{w}(E)}{\mathrm{r}(\partial\mathbf{S})}\right] \\
&= \mathbb{P}\left[\sup_{\mathbf{S}} \|P^{\mathbf{S}}\varepsilon\|^2 - C\mathrm{r}(\partial\mathbf{S}) \log \frac{e \cdot \mathrm{w}(E)}{\mathrm{r}(\partial\mathbf{S})} > 0\right] \\
&= \mathbb{P}\left[\sup_{\mathbf{S}} \exp\left(\eta\|P^{\mathbf{S}}\varepsilon\|^2 - \eta C\mathrm{r}(\partial\mathbf{S}) \log \frac{e \cdot \mathrm{w}(E)}{\mathrm{r}(\partial\mathbf{S})}\right) > 1\right] \\
&\leq \mathbb{E}\left[\sup_{\mathbf{S}} \exp\left(\eta\|P^{\mathbf{S}}\varepsilon\|^2 - \eta C\mathrm{r}(\partial\mathbf{S}) \log \frac{e \cdot \mathrm{w}(E)}{\mathrm{r}(\partial\mathbf{S})}\right)\right],
\end{aligned}$$

where the suprema are over all non-trivial multi-cuts $\mathbf{S}$ of G. Corresponding to the weighting $\mathrm{r} : E \to \mathbb{R}_+$ is a distribution over random spanning trees T of G, satisfying the property (10). Let $\mathbb{E}_\varepsilon$ and $\mathbb{E}_T$ denote the expectations over ε and over T. Then the above yields

$$p_n \leq \mathbb{E}_\varepsilon\left[\sup_{\mathbf{S}} \exp\left(\eta\|P^{\mathbf{S}}\varepsilon\|^2 - \eta C\,\mathbb{E}_T[|T \cap \partial\mathbf{S}|] \log \frac{e \cdot \mathrm{w}(E)}{\mathbb{E}_T[|T \cap \partial\mathbf{S}|]}\right)\right].$$

Observe that for any $c > 0$, $x \mapsto x\log(ec/x)$ is concave over $x \in (0, c)$, so $x \mapsto \exp(-\eta C x \log(ec/x))$ is convex over $x \in (0, c)$. Also, $|T \cap \partial\mathbf{S}| \leq |T| = n - 1 = \mathrm{r}(E) \leq \mathrm{w}(E)$. Then applying Jensen's inequality,

$$\begin{aligned}
p_n &\leq \mathbb{E}_\varepsilon\left[\sup_{\mathbf{S}} \mathbb{E}_T\left[\exp\left(\eta\|P^{\mathbf{S}}\varepsilon\|^2 - \eta C|T \cap \partial\mathbf{S}| \log \frac{e \cdot \mathrm{w}(E)}{|T \cap \partial\mathbf{S}|}\right)\right]\right] \\
&\leq \mathbb{E}_\varepsilon\mathbb{E}_T\left[\sup_{\mathbf{S}} \exp\left(\eta\|P^{\mathbf{S}}\varepsilon\|^2 - \eta C|T \cap \partial\mathbf{S}| \log \frac{e \cdot \mathrm{w}(E)}{|T \cap \partial\mathbf{S}|}\right)\right]. \qquad (16)
\end{aligned}$$

Fix T, and note that $|T \cap \partial\mathbf{S}| \geq 1$ for any non-trivial multi-cut $\mathbf{S}$. We now control the supremum over $\mathbf{S}$ by a union bound over all possible non-empty edge sets $T \cap \partial\mathbf{S}$:

$$\begin{aligned}
&\sup_{\mathbf{S}} \exp\left(\eta\|P^{\mathbf{S}}\varepsilon\|^2 - \eta C|T \cap \partial\mathbf{S}| \log \frac{e \cdot \mathrm{w}(E)}{|T \cap \partial\mathbf{S}|}\right) \\
&\leq \sum_{E' \subseteq T : |E'| \geq 1} \sup_{\mathbf{S} :\, T \cap \partial\mathbf{S} = E'} \exp\left(\eta\|P^{\mathbf{S}}\varepsilon\|^2 - \eta C|E'| \log \frac{e \cdot \mathrm{w}(E)}{|E'|}\right)
\end{aligned}$$

Consider any fixed non-empty subset $E' \subseteq T$, and denote by $\mathbf{S}^* := \mathbf{S}^*(E', T)$ the multi-cut obtained in the following way: Remove the edges E' from T, and let the elements of $\mathbf{S}^*$ be the remaining connected components in the graph with edges $T \setminus E'$. If $\mathbf{S}$ is any other multi-cut satisfying $T \cap \partial\mathbf{S} = E'$, then any two vertices in the same element of $\mathbf{S}^*$ are connected by a path of edges in T that are not cut by $\mathbf{S}$, and hence these vertices belong to the same element of $\mathbf{S}$. Thus $\mathbf{S}^*$ is a refinement of $\mathbf{S}$. Then the range of the projection $P^{\mathbf{S}}$ is a linear subspace of the range of $P^{\mathbf{S}^*}$, and we consequently have $\|P^{\mathbf{S}^*}\varepsilon\|^2 \geq \|P^{\mathbf{S}}\varepsilon\|^2$ (deterministically for any ε). Then

$$p_n \leq \mathbb{E}_T\left[\sum_{E' \subseteq T : |E'| \geq 1} e^{-\eta C|E'| \log \frac{e \cdot \mathrm{w}(E)}{|E'|}} \mathbb{E}_\varepsilon\left[e^{\eta\|P^{\mathbf{S}^*}\varepsilon\|^2}\right]\right].$$

$\mathbf{S}^*$ has $|E'|+1$ elements, so $\|P^{\mathbf{S}^*}\varepsilon\|^2 \sim \chi^2_{|E'|+1}$. Then for $\eta \in (0,1/2)$, the chi-squared moment-generating-function yields $\mathbb{E}_\varepsilon[e^{\eta\|P^{\mathbf{S}^*}\varepsilon\|^2}] = (1-2\eta)^{|E'|+1}$, so

$$p_n \le \mathbb{E}_T\left[\sum_{l=1}^{n-1}\sum_{E'\subseteq T:|E'|=l} e^{-\eta C l\log\frac{e\cdot \mathrm{w}(E)}{l}}(1-2\eta)^{l+1}\right].$$

Applying $|\{E'\subseteq T : |E'|=l\}| = \binom{n-1}{l} \le \exp(l\log\frac{e(n-1)}{l}) \le \exp(l\log\frac{e\cdot\mathrm{w}(E)}{l})$ for every spanning tree T,

$$p_n \le \sum_{l=1}^{n-1} e^{(-\eta C+1)l\log\frac{e\cdot\mathrm{w}(E)}{l}}(1-2\eta)^{l+1} \le \sum_{l=1}^{\lfloor \mathrm{w}(E)\rfloor} e^{(-\eta C+1)l\log\frac{e\cdot\mathrm{w}(E)}{l}}(1-2\eta)^{l+1}.$$

Fixing η to be a constant (say $\eta=1/4$), it is easily verified that for $C>1/\eta$, the function

$$l \mapsto (-\eta C+1)l\log\frac{e\cdot\mathrm{w}(E)}{l} + (l+1)\log(1-2\eta)$$

is convex, so its maximum over $l\in[1,\mathrm{w}(E)]$ is attained at either $l=1$ or $l=\mathrm{w}(E)$. Thus

$$p_n \le \mathrm{w}(E)\max\left((e\cdot\mathrm{w}(E))^{-\eta C+1}(1-2\eta)^2, e^{-(\eta C+1)\mathrm{w}(E)}(1-2\eta)^{\mathrm{w}(E)+1}\right) \le C'\mathrm{w}(E)^{-3}$$

for sufficiently large constants $C, C'>0$, as desired. □

The preceding two lemmas yield risk bounds for $\hat\mu$ in probability via a Cauchy-Schwarz argument.

Lemma A.3. *Let $\delta>0$ and $\tau\ge 0$, and let $\mathrm{w}: E\to\mathbb{R}_+$ be such that $\mathrm{w}\ge \mathrm{r}$ for some $\mathrm{r}\in\mathcal{ST}(G)$.*

(a) For any $\eta>0$, there exist constants $C, C', C''>0$ depending only on η such that for any true signal $\mu_0\in\mathbb{R}^n$, any (fixed) $\mu\in\mathbb{R}^n$, and any $\lambda \ge C\log\mathrm{w}(E)$, with probability at least $1-C'\mathrm{w}(E)^{-3}$ all $(\tau,\delta\mathbb{Z})$-local-minima $\hat\mu$ of (W) satisfy

$$\|\hat\mu-\mu_0\|^2 \le (1+\eta)\|\mu-\mu_0\|^2 + C''\left((\lambda+\tau)\max(\|D\mu\|_{\mathrm{w}},1)+n\delta^2\right).$$

(b) There exist universal constants $C, C', C''>0$ such that for any $s\in[1,\mathrm{w}(E)]$, any true signal $\mu_0\in\mathbb{R}^n$ with $\|D\mu_0\|_{\mathrm{w}}\le s$, and any $\lambda \ge C(1+\log\frac{\mathrm{w}(E)}{s})$, with probability at least $1-C'\mathrm{w}(E)^{-3}$ all $(\tau,\delta\mathbb{Z})$-local-minima $\hat\mu$ of (W) satisfy

$$\|\hat\mu-\mu_0\|^2 \le C''((\lambda+\tau)s+n\delta^2).$$

Proof. For part (a), let $\mu^\delta\in(\delta\mathbb{Z})^n$ be the vector obtained by rounding each entry of μ to the nearest value in $\delta\mathbb{Z}$, and let $\hat\mu$ be any $(\tau,\delta\mathbb{Z})$-local-minimizer. Recall that if the entries of μ^δ take k distinct values, then $\|D\mu^\delta\|_{\mathrm{w}}\ge k-1$ by (12). Then by Lemma A.1,

$$\frac{1}{2}\|Y-\hat\mu\|^2+\lambda\|D\hat\mu\|_{\mathrm{w}} \le \frac{1}{2}\|Y-\mu^\delta\|^2 + 2\lambda\|D\mu^\delta\|_{\mathrm{w}} + \tau(\|D\mu^\delta\|_{\mathrm{w}}+1).$$

Writing $Y=\mu_0+\varepsilon$ and canceling $\|\varepsilon\|^2/2$ from both sides,

$$\frac{1}{2}\|\hat\mu-\mu_0\|^2+\langle\mu_0-\hat\mu,\varepsilon\rangle+\lambda\|D\hat\mu\|_{\mathrm{w}} \le \frac{1}{2}\|\mu^\delta-\mu_0\|^2+\langle\mu_0-\mu^\delta,\varepsilon\rangle+2\lambda\|D\mu^\delta\|_{\mathrm{w}}+\tau(\|D\mu^\delta\|_{\mathrm{w}}+1). \quad (17)$$

Suppose μ^δ induces the multi-cut $\mathbf{S}^\delta$ of G, and $\hat\mu$ induces the multi-cut $\hat{\mathbf{S}}$ of G. Denote by $\mathbf{S}^\delta\vee\hat{\mathbf{S}}$ the least common refinement of $\mathbf{S}^\delta$ and $\hat{\mathbf{S}}$, so

$$\langle\hat\mu-\mu^\delta,\varepsilon\rangle = \langle P^{\mathbf{S}^\delta\vee\hat{\mathbf{S}}}(\hat\mu-\mu^\delta),\varepsilon\rangle = \langle\hat\mu-\mu^\delta, P^{\mathbf{S}^\delta\vee\hat{\mathbf{S}}}\varepsilon\rangle.$$

For any positive constants c_1,c_2 such that $c_1<\frac14$ and $c_1c_2=\frac14$, applying $xy\le c_1x^2+c_2y^2$ and $(x+y)^2\le 2(x^2+y^2)$ yields

$$\begin{aligned}\langle\hat\mu-\mu^\delta,\varepsilon\rangle &\le (\|\hat\mu-\mu_0\|+\|\mu^\delta-\mu_0\|)\|P^{\mathbf{S}^\delta\vee\hat{\mathbf{S}}}\varepsilon\| \\ &\le 2c_1(\|\hat\mu-\mu_0\|^2+\|\mu^\delta-\mu_0\|^2)+c_2\|P^{\mathbf{S}^\delta\vee\hat{\mathbf{S}}}\varepsilon\|^2.\end{aligned}$$

Applying this to (17) and rearranging yields

$$\left(\frac{1}{2}-2c_1\right)\|\hat{\mu}-\mu_0\|^2 \le \left(\frac{1}{2}+2c_1\right)\|\mu^\delta-\mu_0\|^2+c_2\|P^{\mathbf{S}^\delta\vee\hat{\mathbf{S}}}\varepsilon\|^2+2\lambda\|D\mu^\delta\|_{\mathrm{w}}+\tau(\|D\mu^\delta\|_{\mathrm{w}}+1)-\lambda\|D\hat{\mu}\|_{\mathrm{w}}. \tag{18}$$

By Lemma A.2, there exist constants $C, C' > 0$ and an event of probability $1 - C'\mathrm{w}(E)^{-3}$ on which we are guaranteed

$$\|P^{\mathbf{S}^\delta\vee\hat{\mathbf{S}}}\varepsilon\|^2 \le C\max(\mathrm{w}(\partial(\mathbf{S}^\delta\vee\hat{\mathbf{S}})),1)\log\frac{e\cdot\mathrm{w}(E)}{\max(\mathrm{w}(\partial(\mathbf{S}^\delta\vee\hat{\mathbf{S}})),1)}.$$

Note that $\partial(\mathbf{S}^\delta\vee\hat{\mathbf{S}})\subseteq\partial\mathbf{S}^\delta\cup\partial\hat{\mathbf{S}}$, so

$$\mathrm{w}(\partial(\mathbf{S}^\delta\vee\hat{\mathbf{S}})) \le \mathrm{w}(\partial\mathbf{S}^\delta)+\mathrm{w}(\partial\hat{\mathbf{S}}) = \|D\mu^\delta\|_{\mathrm{w}}+\|D\hat{\mu}\|_{\mathrm{w}}. \tag{19}$$

Then on this event,

$$\|P^{\mathbf{S}^\delta\vee\hat{\mathbf{S}}}\varepsilon\|^2 \le C(\max(\|D\mu^\delta\|_{\mathrm{w}},1)+\|D\hat{\mu}\|_{\mathrm{w}})\log(e\cdot\mathrm{w}(E)).$$

If $\lambda \ge c_2C\log(e\cdot\mathrm{w}(E))$, then applying this to (18) yields (on this event)

$$\|\hat{\mu}-\mu_0\|^2 \le \frac{1+4c_1}{1-4c_1}\|\mu^\delta-\mu_0\|^2+C''\left((\lambda+\tau)\max(\|D\mu^\delta\|_{\mathrm{w}},1)\right)$$

for some constant $C'' > 0$. Part (a) follows by noting $\|D\mu^\delta\|_{\mathrm{w}} \le \|D\mu\|_{\mathrm{w}}$,

$$\|\mu^\delta-\mu_0\|^2 \le (1+c_1)\|\mu-\mu_0\|^2+(1+c_1^{-1})\|\mu-\mu^\delta\|^2 \le (1+c_1)\|\mu-\mu_0\|^2+(1+c_1^{-1})n\delta^2,$$

and taking c_1 sufficiently small.

For part (b), let us apply the preceding arguments for $\mu=\mu_0$, so that μ^δ is obtained by rounding the entries of the true signal μ_0. Then $\|\mu^\delta-\mu_0\|^2 \le n\delta^2$, $\|D\mu^\delta\|_{\mathrm{w}} \le \|D\mu_0\|_{\mathrm{w}} \le s$, and (19) implies $\max(\mathrm{w}(\partial(\mathbf{S}^\delta\vee\hat{\mathbf{S}})),1) \le s+\|D\hat{\mu}\|_{\mathrm{w}}$. Observing that $s+\|D\hat{\mu}\|_{\mathrm{w}} \le 2\mathrm{w}(E)$ and that $x\mapsto x\log(ec/x)$ is increasing for $x\in(0,c)$, on an event of probability $1-C'\mathrm{w}(E)^{-3}$ we have

$$\begin{aligned}\|P^{\mathbf{S}^\delta\vee\hat{\mathbf{S}}}\varepsilon\|^2 &\le C\max(\mathrm{w}(\partial(\mathbf{S}^\delta\vee\hat{\mathbf{S}})),1)\log\frac{2e\cdot\mathrm{w}(E)}{\max(\mathrm{w}(\partial(\mathbf{S}^\delta\vee\hat{\mathbf{S}})),1)}\\ &\le C(s+\|D\hat{\mu}\|_{\mathrm{w}})\log\frac{2e\cdot\mathrm{w}(E)}{s+\|D\hat{\mu}\|_w} \le C(s+\|D\hat{\mu}\|_{\mathrm{w}})\log\frac{2e\cdot\mathrm{w}(E)}{s}.\end{aligned}$$

If $\lambda \ge c_2C\log(2e\cdot\mathrm{w}(E)/s)$, then applying this to (18) yields $\|\hat{\mu}-\mu_0\|^2 \le C''((\lambda+\tau)s+n\delta^2)$ for a constant $C'' > 0$, as desired. □

Finally, we obtain bounds in expectation by applying Hölder's inequality and a crude bound on the qth power of the squared-error risk for some $q > 1$.

Lemma A.4. *Let $\delta > 0$, $\tau \ge 0$, and $\hat{\mu}$ be any $(\tau,\delta\mathbb{Z})$-local-minimizer of (W). Then for any $q > 1$, there exists a constant C_q depending only on q such that*

$$\mathbb{E}[\|\hat{\mu}-\mu_0\|^{2q}] \le C_q(\lambda\mathrm{w}(E)+n+n\delta^2+n\tau)^q.$$

Proof. Let $\mu^\delta\in(\delta\mathbb{Z})^n$ denote the vector obtained by rounding each entry of μ_0 to the nearest value in $\delta\mathbb{Z}$. For each vertex $i\in\{1,\dots,n\}$, the vector $\tilde{\mu}$ obtained by replacing $\hat{\mu}_i$ with μ_i^δ and keeping all other coordinates of $\hat{\mu}$ the same is a $\delta\mathbb{Z}$-expansion of $\hat{\mu}$. Since $\hat{\mu}$ is a $(\tau,\delta\mathbb{Z})$-local-minimizer, this implies $F_{\mathrm{w}}(\hat{\mu}) \le F_{\mathrm{w}}(\tilde{\mu})+\tau$ and hence (cancelling terms not depending on the ith coordinate)

$$\frac{1}{2}(\hat{\mu}_i-Y_i)^2 \le \frac{1}{2}(\mu_i^\delta-Y_i)^2+\lambda\sum_{j:\{i,j\}\in E}\mathrm{w}(i,j)+\tau.$$

Summing over $i = 1, \dots, n$ and applying $\|\mu^\delta - Y\|^2 \le 2\|\mu_0 - Y\|^2 + 2\|\mu^\delta - \mu_0\|^2 \le 2\|\mu_0 - Y\|^2 + 2n\delta$,

$$\frac{1}{2}\|\hat{\mu} - Y\|^2 \le \frac{1}{2}\|\mu^\delta - Y\|^2 + 2\lambda \mathrm{w}(E) + n\tau \le \|\mu_0 - Y\|^2 + 2\lambda \mathrm{w}(E) + n\delta^2 + n\tau.$$

Then, applying $(a+b)^q \le 2^{q-1}(a^q + b^q)$,

$$\begin{aligned}\|\hat{\mu} - \mu_0\|^{2q} &\le 2^{2q-1}(\|\hat{\mu} - Y\|^{2q} + \|\mu_0 - Y\|^{2q}) \\ &\le 2^{2q-1}\left(2^{q-1}\left((2\|\mu_0 - Y\|^2)^q + (4\lambda \mathrm{w}(E) + 2n\delta^2 + 2n\tau)^q\right) + \|\mu_0 - Y\|^{2q}\right) \\ &\le C_q(\|\mu_0 - Y\|^{2q} + (\lambda \mathrm{w}(E) + n\delta^2 + n\tau)^q)\end{aligned}$$

for a constant $C_q > 0$. Since $\|\mu_0 - Y\|^2 \sim \chi^2_n$, this implies

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^{2q}] \le C'_q(\lambda \mathrm{w}(E) + n + n\delta^2 + n\tau)^q$$

for a different constant $C'_q > 0$. □

Proof of Theorem 2.14. For $\mu \in \mathbb{R}^n$ and $\lambda, \eta, C'' > 0$, denote

$$M(\mu, \lambda, \eta, C'') := (1+\eta)\|\mu - \mu_0\|^2 + C''\lambda \max(\|D\mu\|_{\mathrm{w}}, 1).$$

Note that $n\delta^2 + \tau \max(\|D\mu\|_{\mathrm{w}}, 1) \le 2\lambda \max(\|D\mu\|_{\mathrm{w}}, 1)$ whenever $\delta \le 1/\sqrt{n}$, $\tau \le 1$, and $\lambda \ge 1$. Applying Hölder's inequality and Lemma A.3(a), there exist constants $C, C', C'' > 0$ such that for any $\mu \in \mathbb{R}^n$, any $\lambda \ge C \log \mathrm{w}(E)$, and any constants $p, q > 1$ such that $\frac{1}{p} + \frac{1}{q} = 1$,

$$\begin{aligned}\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] &\le M(\mu, \lambda, \eta, C'') + \mathbb{E}\left[\|\hat{\mu} - \mu_0\|^2 \mathbb{1}\{\|\hat{\mu} - \mu_0\|^2 > M(\mu, \lambda, \eta, C'')\}\right] \\ &\le M(\mu, \lambda, \eta, C'') + \left(\mathbb{E}[\|\hat{\mu} - \mu_0\|^{2q}]\right)^{1/q}\left(\mathbb{P}[\|\hat{\mu} - \mu_0\|^2 > M(\mu, \lambda, \eta, C'')]\right)^{1/p} \\ &\le M(\mu, \lambda, \eta, C'') + \left(\mathbb{E}[\|\hat{\mu} - \mu_0\|^{2q}]\right)^{1/q}(C' \mathrm{w}(E)^{-3})^{1/p}.\end{aligned}$$

Note that $\mathrm{w}(E) \ge \mathrm{r}(E) = n - 1$, so $\lambda \mathrm{w}(E) \ge n + n\delta^2 + n\tau$ for any $\lambda \ge 3$. Then, choosing $p = 3$, $q = 3/2$, and applying Lemma A.4,

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \le M(\mu, \lambda, \eta, C'') + C'''\lambda$$

for a constant $C''' > 0$. Taking the infimum over $\mu \in \mathbb{R}^n$ yields the desired result. □

Proof of Theorem 2.15. The proof of the upper bound is the same as that of Theorem 2.14, using Lemma A.3(b) instead of Lemma A.3(a): For universal constants $C, C', C'', C''' > 0$, any $\delta \le 1/\sqrt{n}$, $\tau \le 1$, $\lambda \ge C(1 + \log \frac{\mathrm{w}(E)}{s})$, and $p = 3$ and $q = 3/2$,

$$\mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \le C''\lambda s + \left(\mathbb{E}[\|\hat{\mu} - \mu_0\|^{2q}]\right)^{1/q}(C' \mathrm{w}(E)^{-3})^{1/p} \le C'''\lambda s.$$

For the lower bound, denote $\mathrm{w}(i) = \sum_{j:\{i,j\}\in E} \mathrm{w}(i,j)$, and consider the vertex subset $V' = \{i : \mathrm{w}(i) \le 4\mathrm{w}(E)/n\}$. Then the class $\{\mu_0 \in \mathbb{R}^n : \|D\mu_0\|_{\mathrm{w}} \le s\}$ contains the class of sparse vectors

$$B_0 := \left\{\mu_0 \in \mathbb{R}^n : \mathrm{supp}(\mu_0) \subseteq V',\ \|\mu_0\|_0 \le \left\lfloor \frac{sn}{4\mathrm{w}(E)} \right\rfloor \right\},$$

where $\mathrm{supp}(\mu_0)$ denotes the set of vertices i for which $\mu_{0,i} \ne 0$. (This is because for any $\mu_0 \in B_0$, $\|D\mu_0\|_{\mathrm{w}} \le \sum_{i\in\mathrm{supp}(\mu_0)} \mathrm{w}(i) \le \lfloor \frac{sn}{4\mathrm{w}(E)} \rfloor \frac{4\mathrm{w}(E)}{n} \le s$.) Then

$$\inf_{\hat{\mu}} \sup_{\mu_0 : \|D\mu_0\|_{\mathrm{w}} \le s} \mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \ge \inf_{\hat{\mu}} \sup_{\mu_0 \in B_0} \mathbb{E}[\|\hat{\mu} - \mu_0\|^2].$$

For $s \ge 4\mathrm{w}(E)/n$, the lower-bound for the sparse normal-means problem (see e.g. Theorem 1(b) of [RWY11] in the case $X = \sqrt{n}\,\mathrm{Id}$) yields, for a universal constant $c > 0$,

$$\inf_{\hat{\mu}} \sup_{\mu_0 \in B_0} \mathbb{E}[\|\hat{\mu} - \mu_0\|^2] \ge c \left\lfloor \frac{sn}{4\mathrm{w}(E)} \right\rfloor \log \frac{|V'|}{\lfloor \frac{sn}{4\mathrm{w}(E)} \rfloor} \ge c \frac{sn}{8\mathrm{w}(E)} \log \frac{4\mathrm{w}(E)|V'|}{sn}.$$

Note that $\mathrm{w}(i) > 4\mathrm{w}(E)/n$ for all $i \notin V'$, so

$$2\mathrm{w}(E) = \sum_{i=1}^{n} \mathrm{w}(i) > \frac{4\mathrm{w}(E)}{n}(n - |V'|).$$

Then $|V'| > n/2$, and the lower bound follows. □

Proof of Theorem 2.19. Let $\mu^\delta \in (\delta\mathbb{Z})^n$ denote the (constant) vector obtained by rounding the value in μ_0 to the closest value in $\delta\mathbb{Z}$. As μ^δ is a $\delta\mathbb{Z}$-expansion of $\hat{\mu}$, and $\hat{\mu}$ is a $(\tau, \delta\mathbb{Z})$-local-minimizer, we have

$$F_{\mathrm{w}}(\hat{\mu}) = \frac{1}{2}\|Y - \hat{\mu}\|^2 + \lambda\|D\hat{\mu}\|_{\mathrm{w}} \leq \frac{1}{2}\|Y - \mu^\delta\|^2 + \tau = F_{\mathrm{w}}(\mu^\delta) + \tau.$$

Denoting by $\mathbf{S}$ the multi-cut induced by $\hat{\mu}$, the same steps as leading to (18) yield, for any positive constants $c_1 < \frac{1}{4}$ and $c_2 > 1$ such that $c_1 c_2 = \frac{1}{4}$,

$$\left(\frac{1}{2} - 2c_1\right)\|\hat{\mu} - \mu_0\|^2 \leq \left(\frac{1}{2} + 2c_1\right) n\delta^2 + c_2\|P^{\mathbf{S}}\varepsilon\|^2 + \tau - \lambda\|D\hat{\mu}\|_{\mathrm{w}}.$$

Hence, for $\delta \leq 1/\sqrt{n}$ and $\tau \leq 1$,

$$\lambda\|D\hat{\mu}\|_{\mathrm{w}} \leq C''(1 + \|P^{\mathbf{S}}\varepsilon\|^2)$$

for a constant $C'' > 0$. Lemma A.2 implies, with probability at least $1 - C'\mathrm{w}(E)^{-3} \geq 1 - C'(n-1)^{-3}$,

$$\|P^{\mathbf{S}}\varepsilon\|^2 \leq C\max(\mathrm{w}(\partial\mathbf{S}), 1)\log(e \cdot \mathrm{w}(E)) = C\max(\|D\hat{\mu}\|_{\mathrm{w}}, 1)\log(e \cdot \mathrm{w}(E)).$$

Recall by (12) that $\|D\hat{\mu}\|_{\mathrm{w}} \geq 1$ if $\|D\hat{\mu}\|_{\mathrm{w}} \neq 0$. Thus, if $\lambda > C''(1 + C\log(e \cdot \mathrm{w}(E)))$, then the above two statements imply $\|D\hat{\mu}\|_{\mathrm{w}} = 0$, and hence $\hat{\mu}$ is constant with probability at least $1 - C'(n-1)^{-3}$, as desired. □

Appendix B. l_1/total-variation lower bound for the linear chain

In this appendix, we prove Theorems 2.9 and 2.10. By rescaling, we will assume $\sigma^2 = 1$. We say that k consecutive intervals $\mathbf{S} := (S_1, \ldots, S_k)$ partition $(1, \ldots, n)$ if there are integers $1 = i_1 < i_2 < \ldots < i_k < i_{k+1} = n + 1$ such that $S_\alpha = (i_\alpha, i_\alpha + 1, \ldots, i_{\alpha+1} - 1)$ for each α. We associate to each such partition $\mathbf{S}$ two orthogonal projections $P^{\mathbf{S}} : \mathbb{R}^n \to \mathbb{R}^n$ and $Q^{\mathbf{S}} : \mathbb{R}^n \to \mathbb{R}^n$, such that $P^{\mathbf{S}}$ projects onto the k-dimensional subspace of vectors assuming a constant value on each S_α, and $Q^{\mathbf{S}}$ projects onto the $(n-k)$-dimensional orthogonal complement of this subspace. More formally, for each α and each $i \in S_\alpha$, $P^{\mathbf{S}}$ is defined by $(P^{\mathbf{S}}\mu)_i = |S_\alpha|^{-1}\sum_{j \in S_\alpha}\mu_j$, and $Q^{\mathbf{S}} = \mathrm{Id} - P^{\mathbf{S}}$. We denote by $P^{\mathbf{S}}(\mathbb{R}^n)$ the range of $P^{\mathbf{S}}$, and for any $v \in P^{\mathbf{S}}(\mathbb{R}^n)$ we denote by v_{S_α} the (constant) value of v over S_α.

We will say that $\mu \in \mathbb{R}^n$ induces such a partition $\mathbf{S} = (S_1, \ldots, S_k)$ and a sign vector $\mathbf{s} = (s_1, \ldots, s_{k-1}) \in \{-1, 1\}^{k-1}$ if $\mu \in P^{\mathbf{S}}(\mathbb{R}^n)$ and $\mathrm{sign}(\mu_{S_{\alpha+1}} - \mu_{S_\alpha}) = s_\alpha$ for each $\alpha = 1, \ldots, k-1$. Recall the vertex-edge incidence matrix $D \in \mathbb{R}^{n \times (n-1)}$: In this appendix, for the linear chain graph, let us fix the sign convention for D so that $D\mu = (\mu_2 - \mu_1, \mu_3 - \mu_2, \ldots, \mu_n - \mu_{n-1})$.

The following lemma is an implication of the subgradient condition for $\hat{\mu}$ minimizing (TV); a similar result was stated as Lemmas 2.1 and A.1 of [Rin09].

Lemma B.1. *Let $\lambda > 0$. Fix k consecutive intervals $\mathbf{S} := (S_1, \ldots, S_k)$ that partition $(1, \ldots, n)$, and fix a sign vector $\mathbf{s} := (s_1, \ldots, s_{k-1}) \in \{-1, 1\}^{k-1}$. Define $v \in P^{\mathbf{S}}(\mathbb{R}^n)$ such that, for each $\alpha = 1, \ldots, k$,*

$$v_{S_\alpha} = \frac{\lambda(s_\alpha - s_{\alpha-1})}{|S_\alpha|}, \tag{20}$$

where we set $s_0 = s_k := 0$. Define also

$$W := \{w \in [-1, 1]^{n-1} : w_i = s_\alpha \text{ if } i \in S_\alpha \text{ and } i + 1 \in S_{\alpha+1} \text{ for some } \alpha \in \{1, \ldots, k-1\}\}, \tag{21}$$

$$K := \{x \in P^{\mathbf{S}}(\mathbb{R}^n) : \operatorname{sign}(x_{S_{\alpha+1}} - x_{S_\alpha}) = s_\alpha \text{ for all } \alpha = 1, \dots, k-1\}. \tag{22}$$

Then the following two events are equivalent:

(1) The minimizer $\hat{\mu}$ of (TV) induces the partition $\mathbf{S}$ and the signs $\mathbf{s}$.
(2) $P^{\mathbf{S}}Y + v \in K$ and $Q^{\mathbf{S}}Y - v \in \lambda D^T W$.

Furthermore, if this event holds, then $\hat{\mu} = P^{\mathbf{S}}Y + v$.

Proof. Suppose $\hat{\mu}$ minimizes (TV) and induces $\mathbf{S}$ and $\mathbf{s}$. Then W is precisely the subdifferential of the l_1-norm at $D\hat{\mu}$, so $\lambda D^T W$ is the subdifferential of $\lambda\|D\hat{\mu}\|_1$ with respect to $\hat{\mu}$. Since $\hat{\mu}$ minimizes (TV), it satisfies the subgradient condition $Y - \hat{\mu} \in \lambda D^T W$. Apply $P^{\mathbf{S}}$ to both sides, noting that $P^{\mathbf{S}}\hat{\mu} = \hat{\mu}$ and $P^{\mathbf{S}}x = -v$ for any $x \in \lambda D^T W$. Then $\hat{\mu} = P^{\mathbf{S}}Y + v$. Since $\hat{\mu}$ induces the signs $\mathbf{s}$, we have $P^{\mathbf{S}}Y + v \in K$. Also, $Y - \hat{\mu} = Q^{\mathbf{S}}Y - v \in \lambda D^T W$, as desired.

Conversely, suppose $P^{\mathbf{S}}Y + v \in K$ and $Q^{\mathbf{S}}Y - v \in \lambda D^T W$. Defining $\hat{\mu} = P^{\mathbf{S}}Y + v$, the first condition implies that $\hat{\mu}$ induces $\mathbf{S}$ and $\mathbf{s}$, while the second implies $Y - \hat{\mu} \in \lambda D^T W$. Since $\lambda D^T W$ is the subdifferential of $\lambda\|D\hat{\mu}\|_1$ with respect to $\hat{\mu}$, this implies that $\hat{\mu}$ minimizes (TV). (The objective (TV) is strictly convex, so $\hat{\mu}$ is the unique minimizer.) □

Using this characterization, we may lower-bound the number of intervals in the partition induced by $\hat{\mu}$ when the true signal μ_0 is 0:

Proof of Theorem 2.9(a). For a fixed partition $\mathbf{S} := (S_1, \dots, S_k)$ into k consecutive intervals, let us first bound

$$p_{\mathbf{S}} := \mathbb{P}\left[Q^{\mathbf{S}}Y \in \lambda Q^{\mathbf{S}} D^T [-1, 1]^{n-1}\right].$$

Denoting by B_m the Euclidean unit ball in $\mathbb{R}^m$ centered at 0, clearly

$$p_{\mathbf{S}} \le \mathbb{P}[Q^{\mathbf{S}}Y \in \lambda\sqrt{n} Q^{\mathbf{S}} D^T \mathsf{B}_{n-1}].$$

The range of D^T is the $(n-1)$-dimensional space orthogonal to the all-1's vector, and $Q^{\mathbf{S}}$ is a projection onto an $(n-k)$-dimensional subspace of this range—hence $Q^{\mathbf{S}}D^T$ has rank $n-k$. Let $Q^{\mathbf{S}}D^T = U\Gamma V^T$ denote the (reduced) singular value decomposition of $Q^{\mathbf{S}}D^T$, where $\Gamma = \operatorname{diag}(\gamma_1, \dots, \gamma_{n-k})$ contains the singular values $\gamma_1 \le \dots \le \gamma_{n-k}$, and $U \in \mathbb{R}^{n\times(n-k)}$ and $V \in \mathbb{R}^{(n-1)\times(n-k)}$ have orthonormal columns. Then, as $V^T\mathsf{B}_{n-1} = \mathsf{B}_{n-k}$, we have

$$p_{\mathbf{S}} \le \mathbb{P}[U^T Q^{\mathbf{S}} Y \in \lambda\sqrt{n}\Gamma\mathsf{B}_{n-k}].$$

The set $\mathcal{E} := \lambda\sqrt{n}\Gamma\mathsf{B}_{n-k}$ is an ellipsoid in $\mathbb{R}^{n-k}$, whose principal axes are aligned with the standard basis and have lengths $\lambda\sqrt{n}\gamma_i$. The vector $\varepsilon' := U^T Q^{\mathbf{S}} Y$ of length $n-k$ has i.i.d. entries distributed as Normal$(0,1)$ (when $\mu_0 = 0$.) Then, for any $l \in \{1, \dots, n-k\}$, letting $\varepsilon'_{1:l}$ denote the first l coordinates of ε' and $\mathcal{E}_{1:l}$ denote the projection of $\mathcal{E}$ onto the first l dimensions,

$$p_{\mathbf{S}} \le \mathbb{P}[\varepsilon'_{1:l} \in \mathcal{E}_{1:l}] \le (2\pi)^{-l/2} \operatorname{Vol}(\mathcal{E}_{1:l}) = \frac{\left(\lambda\sqrt{\frac{n}{2}}\right)^l}{\Gamma(\frac{l}{2}+1)} \prod_{i=1}^{l} \gamma_i.$$

Here, $\operatorname{Vol}(\cdot)$ denotes the volume in $\mathbb{R}^l$, and we have used $\operatorname{Vol}(\mathsf{B}_l) = \pi^{l/2}/\Gamma(\frac{l}{2}+1)$. Note, by Cauchy interlacing, that $\gamma_i \le s_{i+k-1}(D^T)$ where $s_{i+k-1}(D^T)$ denotes the $i+k-1^{\text{th}}$ smallest non-zero singular value of D^T. (Equality holds if the $k-1$ directions in the column span of D^T that are projected out by $Q^{\mathbf{S}}$ are exactly those corresponding to the smallest $k-1$ non-zero singular values of D^T.) As $D^T D$ is the Laplacian of the linear chain graph, which has non-zero eigenvalues $4\sin^2(i\pi/2n) \le (i\pi/n)^2$ for $i = 1, \dots, n-1$ [AJM85], this implies $\gamma_i \le (i+k-1)\pi/n$ and hence, for any $l \in \{1, \dots, n-k\}$,

$$p_{\mathbf{S}} \le \frac{\left(\frac{\lambda\pi}{\sqrt{2n}}\right)^l}{\Gamma(\frac{l}{2}+1)} \frac{(l+k-1)!}{(k-1)!} = \left(\frac{\lambda\pi}{\sqrt{2n}}\right)^l \frac{l!}{\Gamma(\frac{l}{2}+1)} \binom{l+k-1}{k-1}.$$

Applying the Gamma function duplication formula $\sqrt{\pi} l! = \Gamma(\frac{l}{2}+1)\Gamma(\frac{l}{2}+\frac{1}{2})2^l$, and $\binom{l+k-1}{k-1} \leq \binom{n-1}{k-1} \leq \binom{n}{k}$, this yields

$$p_{\mathbf{S}} \leq \left(\lambda\pi\sqrt{2/n}\right)^l \frac{\Gamma(\frac{l+1}{2})}{\sqrt{\pi}} \binom{n}{k}. \tag{23}$$

Fix $k \geq 1$. Suppose $\hat{\mu}$ induces the partition $\hat{\mathbf{S}}$ having $\hat{k}$ elements, and $\hat{k} \leq k$. Then Lemma B.1 implies $Q^{\hat{\mathbf{S}}} Y - v \in \lambda D^T W$, where v and W are defined in terms of $\hat{\mathbf{S}}$ and $\hat{\mathbf{s}}$. Applying $Q^{\hat{\mathbf{S}}}$ to both sides and noting $Q^{\hat{\mathbf{S}}} v = 0$,

$$Q^{\hat{\mathbf{S}}} Y \in \lambda Q^{\hat{\mathbf{S}}} D^T W \subseteq \lambda Q^{\hat{\mathbf{S}}} D^T [-1,1]^{n-1}.$$

Thus there exists some partition $\mathbf{S}$ into exactly k consecutive intervals (equal to or refining $\hat{\mathbf{S}}$) for which applying $Q^{\mathbf{S}}$ to both sides above yields $Q^{\mathbf{S}} Y \in \lambda Q^{\mathbf{S}} D^T [-1,1]^{n-1}$. As there are $\binom{n-1}{k-1} \leq \binom{n}{k}$ such partitions, applying a union bound, (23), and $\Gamma(x+1) \leq \sqrt{2\pi}((x+0.5)/e)^{x+0.5}$ for any $x > 0$ (see Theorem 1.5 of [Bat08]),

$$\begin{aligned}
\mathbb{P}[\hat{k} \leq k] &\leq \binom{n}{k}^2 \left(\lambda\pi\sqrt{2/n}\right)^l \frac{\Gamma(\frac{l+1}{2})}{\sqrt{\pi}} \\
&\leq \sqrt{2} \exp\left(2k \log\frac{en}{k} + l \log \lambda\pi\sqrt{2/n} + \frac{l}{2}\log\frac{l}{2} - \frac{l}{2}\right).
\end{aligned}$$

Taking $l = \lfloor \frac{n}{\max(\lambda^2,1)\pi^2} \rfloor$ ensures $l \log \lambda\pi\sqrt{2/n} + \frac{l}{2}\log\frac{l}{2} \leq 0$, and taking $k = \lfloor \frac{cn}{\max(\lambda^2,1)\log n} \rfloor$ for a sufficiently small constant $c > 0$ then ensures $\mathbb{P}[\hat{k} \leq k] \leq Ce^{-c'n/\max(\lambda^2,1)}$, for constants $C, c' > 0$, as desired. □

To show part (b) of Theorem 2.9, we will argue that on the event where $\hat{\mu}$ induces a partition into at least k intervals, the squared-error of $\hat{\mu}$ is typically also at least k (up to logarithmic factors). We establish this using the next two lemmas.

Suppose that $\hat{\mu}$ induces the partition $\mathbf{S} = (S_1, \ldots, S_k)$ and the signs $\mathbf{s} = (s_1, \ldots, s_{k-1}) \in \{-1,1\}^{k-1}$. For any $\alpha \in \{1, \ldots, k\}$, let us call S_α a local maximum interval if $s_{\alpha-1} \geq 0$ and $s_\alpha \leq 0$, and a local minimum interval if $s_{\alpha-1} \leq 0$ and $s_\alpha \geq 0$, where by convention $s_0 = s_k := 0$.

Lemma B.2. *Let $\lambda > 0$, let $\hat{\mu}$ be the minimizer of (TV), and suppose μ_0 is such that $\|\mu_0\|_\infty \leq M$. There exists a constant $c > 0$ such that with probability at least $1 - 2/n$, every local minimum interval and local maximum interval of the partition and signs induced by $\hat{\mu}$ has length at least $c \min(\frac{\lambda^2}{\log n}, \frac{\lambda}{M})$.*

Proof. For any $v \in \mathbb{R}^n$ and any interval $S \subseteq (1, \ldots, n)$, denote $\bar{v}_S = |S^{-1}| \sum_{i \in S} v_i$. Define the event

$$\mathcal{E} := \left\{\text{there exists an interval } S \subseteq (1, \ldots, n) : |\bar{\varepsilon}_S| > \sqrt{(6 \log n)/|S|}\right\}.$$

By a Gaussian tail bound and union bound, $\mathbb{P}[\mathcal{E}] \leq 2n^2 e^{-3\log n} = 2/n$.

Let $\mathbf{S} = (S_1, \ldots, S_k)$ be the partition into consecutive intervals induced by $\hat{\mu}$. Define

$$A := \{\alpha : S_\alpha \text{ is a local maximum interval or local minimum interval}\}.$$

Consider $\alpha \in A$ corresponding to the smallest length $|S_\alpha|$. Let $\beta \in A$ be the element immediately before or after α in A, and note that if S_α is a local maximum, then S_β is a local minimum, and vice versa. Then, letting $v \in P^{\mathbf{S}}(\mathbb{R}^n)$ be as defined in (20), $|\bar{Y}_{S_\alpha} - \bar{Y}_{S_\beta}|$ exceeds $|\hat{\mu}_{S_\alpha} - \hat{\mu}_{S_\beta}|$ by $|v_{S_\alpha}| + |v_{S_\beta}|$, so we must have

$$|\bar{Y}_{S_\alpha} - \bar{Y}_{S_\beta}| \geq |v_{S_\alpha}| + |v_{S_\beta}| \geq \frac{\lambda}{|S_\alpha|} + \frac{\lambda}{|S_\beta|} \geq \frac{\lambda}{|S_\alpha|}.$$

On the event $\mathcal{E}$,

$$|\bar{Y}_{S_\alpha} - \bar{Y}_{S_\beta}| \le |\bar{\varepsilon}_{S_\alpha} - \bar{\varepsilon}_{S_\beta}| + 2M \le \sqrt{\frac{6\log n}{|S_\alpha|}} + \sqrt{\frac{6\log n}{|S_\beta|}} + 2M \le 2\sqrt{\frac{6\log n}{|S_\alpha|}} + 2M,$$

where the last bound recalls that S_α has smallest length among $\alpha \in A$. Then on $\mathcal{E}$,

$$0 \le 2M|S_\alpha| + 2\sqrt{6(\log n)|S_\alpha|} - \lambda,$$

which implies

$$\sqrt{|S_\alpha|} \ge \frac{\sqrt{24\log n + 8M\lambda} - 2\sqrt{6\log n}}{4M} = \frac{2\lambda}{\sqrt{24\log n + 8M\lambda} + 2\sqrt{6\log n}} \ge c\min\left(\frac{\lambda}{\sqrt{\log n}}, \sqrt{\frac{\lambda}{M}}\right).$$

□

Lemma B.3. *Fix any $\mu \in \mathbb{R}^k$. Let $\xi \in \mathbb{R}^k$ be a random vector distributed as $\xi \sim \mathrm{Normal}(\mu, \mathrm{Id})$, and let $K \subseteq \mathbb{R}^k$ be any convex cone with vertex 0 and non-empty interior. Then there exists a universal constant $c > 0$ such that*

$$\mathbb{E}[\|\xi\|^2 \mid \xi \in K] \ge \frac{ck^2}{\max(\|\mu\|^2, k)}.$$

Proof. Consider a change of variables to $S = \|\xi\|$ and Θ a set of $k-1$ angular variables representing the angle of ξ. Since K is a cone with vertex 0, the condition $\xi \in K$ is a function only of Θ. Hence the joint density of S and Θ conditional on $\xi \in K$ is given by

$$p(s,\theta) = c(k,\theta,K)s^{k-1}e^{-\frac{\|\xi(s,\theta)-\mu\|^2}{2}},$$

for some quantity $c(k,\theta,K)$ independent of s. Note that for any $s \ge t \ge 0$ and any θ,

$$\begin{aligned}\|\xi(t,\theta)-\mu\|^2 - \|\xi(s,\theta)-\mu\|^2 &= t^2 - s^2 + 2\langle \xi(s,\theta) - \xi(t,\theta), \mu\rangle \\ &\ge t^2 - s^2 - 2(s-t)\|\mu\| = (t+\|\mu\|)^2 - (s+\|\mu\|)^2.\end{aligned}$$

Letting R denote a positive random variable with density function

$$q(r) = I(\mu)r^{k-1}e^{-\frac{(r+\|\mu\|)^2}{2}},$$

where $I(\mu)$ is the normalization constant such that $\int_0^\infty q(r)dr = 1$, this implies $p(s,\theta) \ge \frac{q(s)}{q(t)}p(t,\theta)$. Integrating over θ and denoting by $p(s)$ the marginal distribution of S conditional on $\xi \in K$, this implies $p(s) \ge \frac{q(s)}{q(t)}p(t)$. The inequality is reversed for $0 \le s \le t$, so

$$\begin{aligned}\mathbb{P}[\|\xi\| \ge t \mid \xi \in K] = \frac{\int_t^\infty p(s)ds}{\int_0^t p(s)ds + \int_t^\infty p(s)ds} &\ge \frac{\frac{p(t)}{q(t)}\int_t^\infty q(s)ds}{\int_0^t p(s)ds + \frac{p(t)}{q(t)}\int_t^\infty q(s)ds} \\ &\ge \frac{\frac{p(t)}{q(t)}\int_t^\infty q(s)ds}{\frac{p(t)}{q(t)}\int_0^t q(s)ds + \frac{p(t)}{q(t)}\int_t^\infty q(s)ds} = \mathbb{P}[R \ge t].\end{aligned}$$

That is, the distribution of $\|\xi\|$ conditional on $\xi \in K$ stochastically dominates the (unconditional) distribution of R.

Thus, to conclude the proof of the lemma, it suffices to lower bound

$$\mathbb{E}[R^2] = \int_0^\infty r^{k+1}e^{-\frac{(r+\|\mu\|)^2}{2}}dr \Big/ \int_0^\infty r^{k-1}e^{-\frac{(r+\|\mu\|)^2}{2}}dr =: M_{k+1}/M_{k-1},$$

where M_l is the l^{th} moment of the Normal$(-\|\mu\|, 1)$ distribution truncated on the left at 0. Integration by parts yields the recurrence $M_{k+1} = -\|\mu\| M_k + k M_{k-1}$, and Cauchy-Schwarz yields $M_{k+1} M_{k-1} \geq M_k^2$. Then

$$\frac{M_{k+1}}{M_{k-1}} + \|\mu\| \sqrt{\frac{M_{k+1}}{M_{k-1}}} - k \geq \frac{M_{k+1}}{M_{k-1}} + \|\mu\| \frac{M_k}{M_{k-1}} - k \geq 0,$$

so

$$\sqrt{\frac{M_{k+1}}{M_{k-1}}} \geq \frac{\sqrt{\|\mu\|^2 + 4k} - \|\mu\|}{2} = \frac{2k}{\|\mu\| + \sqrt{\|\mu\|^2 + 4k}},$$

and the result follows by taking the square. □

Proof of Theorem 2.9(b). For any positive constant n_0, it suffices to consider graphs with $n \geq n_0$. (For $n < n_0$, the lower bound trivially holds by adjusting c' based on n_0, as the risk is clearly non-zero for any graph.)

Let $\hat{\mu}$ induce the partition $\hat{\mathbf{S}}$ and signs $\hat{\mathbf{s}}$. For any fixed partition $\mathbf{S} = (S_1, \ldots, S_k)$ and signs $\mathbf{s}$, by Lemma B.1,

$$\mathbb{E}\big[\|\hat{\mu}\|^2 \mid \hat{\mathbf{S}} = \mathbf{S}, \hat{\mathbf{s}} = \mathbf{s}\big] = \mathbb{E}\big[\|P^{\mathbf{S}} Y + v\|^2 \mid P^{\mathbf{S}} Y + v \in K,\, Q^{\mathbf{S}} Y - v \in \lambda D^T W\big],$$

where v, K, and W are fixed and defined by $\mathbf{S}$ and $\mathbf{s}$. Since the projections $P^{\mathbf{S}}$ and $Q^{\mathbf{S}}$ are orthogonal to each other, $P^{\mathbf{S}} Y$ is independent of $Q^{\mathbf{S}} Y$, so we may drop the conditioning on the event $Q^{\mathbf{S}} Y - v \in \lambda D^T W$. Define $\xi \in \mathbb{R}^k$ by $\xi_\alpha = |S_\alpha|^{1/2} (\bar{Y}_{S_\alpha} + v_{S_\alpha})$ for each $\alpha = 1, \ldots, k$, where $\bar{Y}_{S_\alpha} = |S_\alpha|^{-1} \sum_{i \in S_\alpha} Y_i$. Then $\|P^{\mathbf{S}} Y + v\|^2 = \|\xi\|^2$, $\xi \sim \text{Normal}(w, \text{Id})$ for $w = (|S_1|^{1/2} v_{S_1}, \ldots, |S_k|^{1/2} v_{S_k})$, and the condition $P^{\mathbf{S}} Y + v \in K$ is equivalent to $\xi \in \tilde{K}$ for some convex cone $\tilde{K} \in \mathbb{R}^k$ with vertex 0 and non-empty interior. Applying Lemma B.3, for a constant $c' > 0$,

$$\mathbb{E}\big[\|\hat{\mu}\|^2 \mid \hat{\mathbf{S}} = \mathbf{S}, \hat{\mathbf{s}} = \mathbf{s}\big] = \mathbb{E}[\|\xi\|^2 \mid \xi \in \tilde{K}] \geq \frac{c' k^2}{\max(\|w\|^2, k)}. \tag{24}$$

Let $\hat{k}$, $\hat{v}$, and $\hat{w}$ denote k, v, and w as defined above for the random induced partition and signs $\hat{\mathbf{S}}$ and $\hat{\mathbf{s}}$. By part (a) of this theorem (already established), for some constant n_0 and any $n \geq n_0$, on an event of probability at least 3/4, $\hat{k} \geq \frac{c' n}{\max(\lambda^2, 1) \log n}$. Note that $\hat{v}$ is only non-zero on those intervals $\hat{S}_\alpha$ which are local maxima or local minima, and that the (constant) value of $\hat{v}$ on any such interval has magnitude at most $2\lambda / |\hat{S}_\alpha|$. Applying Lemma B.2 with $M = 0$, on a different event of probability at least 3/4, every local minimum and local maximum interval of $\hat{\mathbf{S}}$ has length $|\hat{S}_\alpha| \geq c' \max(\lambda^2, 1)/(\log n)$. (For $\lambda < 1$ this statement is trivial.) Then, on this event, there are at most $\frac{n \log n}{c' \max(\lambda^2, 1)}$ such intervals, and hence

$$\|\hat{w}\|^2 = \sum_{\alpha=1}^{\hat{k}} |\hat{S}_\alpha| \hat{v}_{\hat{S}_\alpha}^2 \leq \sum_{\alpha : \hat{S}_\alpha \text{ local max or min}} \frac{4\lambda^2}{|\hat{S}_\alpha|} \leq \frac{n \log n}{c' \max(\lambda^2, 1)} \frac{4 \log n}{c'} \leq \frac{4n (\log n)^2}{c'^2 \max(\lambda^2, 1)}.$$

Taking the full expectation of (24) over the intersection of these two events of probability at least 1/2, the theorem follows. □

Finally, we prove Theorem 2.10 using Theorem 2.9(b) and Lemma B.2:

Proof of Theorem 2.10. Let c_0 denote the constant in Lemma B.2. If

$$\lambda^2 \leq \frac{3 \log n}{c_0} \sqrt{\frac{n\Delta}{s}},$$

then the lower bound follows from Theorem 2.9(b) by considering the risk at $\mu_0 = 0$. Otherwise, construct a signal μ_0 by the following procedure:

(1) Choose
$$k = \Big\lfloor \min\Big(\frac{s}{2}, \frac{n}{3\Delta}\Big)\Big\rfloor$$
disjoint intervals $I_1, \ldots, I_k$ in $(1, \ldots, n)$, each of length 3Δ.
(2) Set $M = c_0\lambda/(3\Delta)$. For each $\alpha = 1, \ldots, k$, divide I_α into three sub-intervals of length Δ, and set $\mu_{0,i} = M$ for each i belonging to the middle such sub-interval.
(3) Set $\mu_{0,i} = 0$ for all remaining indices.

This signal satisfies $\|D\mu_0\|_0 \le 2k \le s$ and $\Delta(\mu_0) \ge \Delta$. We have
$$\frac{\lambda^2}{\log n} > \frac{3}{c_0}\sqrt{\frac{n\Delta}{s}} > \frac{3\Delta}{c_0} = \frac{\lambda}{M},$$
so Lemma B.2 implies that on an event $\mathcal{E}$ of probability at least $1 - 2/n$, every local maximum or local minimum interval for the partition and signs induced by $\hat{\mu}$ has length at least 3Δ. On this event $\mathcal{E}$, there cannot be a sub-interval of any I_α that is a local maximum or local minimum, so the estimate $\hat{\mu}$ must be monotonically increasing or monotonically decreasing over each interval I_α. Then it is easily verified that
$$\|\hat{\mu} - \mu_0\|^2 \ge \sum_{\alpha=1}^{k}\sum_{i\in I_\alpha} \|\hat{\mu}_i - \mu_{0,i}\|^2 \ge ck\Delta M^2 \ge c'(\log n)\sqrt{\frac{sn}{\Delta}}$$
for constants $c, c' > 0$, where the last inequality uses $\lambda^2 \gtrsim (\log n)\sqrt{n\Delta/s}$ and $k \gtrsim s$. The result follows upon taking the expectation over $\mathcal{E}$. □

References

[ABBDL10] Louigi Addario-Berry, Nicolas Broutin, Luc Devroye, and Gábor Lugosi. On combinatorial testing problems. *The Annals of Statistics*, 38(5):3063–3092, 2010.

[ACCD11] Ery Arias-Castro, Emmanuel J Candès, and Arnaud Durand. Detection of an anomalous cluster in a network. *The Annals of Statistics*, 39(1):278–304, 2011.

[ACCHZ08] Ery Arias-Castro, Emmanuel J Candès, Hannes Helgason, and Ofer Zeitouni. Searching for a trail of evidence in a maze. *The Annals of Statistics*, 36(4):1726–1757, 2008.

[ACDH05] Ery Arias-Castro, David L Donoho, and Xiaoming Huo. Near-optimal detection of geometric objects by fast multiscale methods. *IEEE Transactions on Information Theory*, 51(7):2402–2425, 2005.

[ACG13] Ery Arias-Castro and Geoffrey R Grimmett. Cluster detection in networks using percolation. *Bernoulli*, 19(2):676–719, 2013.

[AG02] Elena Andreou and Eric Ghysels. Detecting multiple breaks in financial market volatility dynamics. *Journal of Applied Econometrics*, 17(5):579–600, 2002.

[AIL99] Reena Aggarwal, Carla Inclan, and Ricardo Leal. Volatility in emerging stock markets. *Journal of Financial and Quantitative Analysis*, 34(1):33–55, 1999.

[AJM85] William N Anderson Jr and Thomas D Morley. Eigenvalues of the Laplacian of a graph. *Linear and multilinear algebra*, 18(2):141–145, 1985.

[AL89] Ivan E Auger and Charles E Lawrence. Algorithms for the optimal identification of segment neighborhoods. *Bulletin of mathematical biology*, 51(1):39–54, 1989.

[Bat08] Necdet Batir. Inequalities for the gamma function. *Archiv der Mathematik*, 91(6):554–563, 2008.

[BBM99] Andrew Barron, Lucien Birgé, and Pascal Massart. Risk bounds for model selection via penalization. *Probability Theory and Related Fields*, 113(3):301–413, 1999.

[BBM00] Jerome V Braun, RK Braun, and H-G Muller. Multiple changepoint fitting via quasilikelihood, with application to DNA sequence segmentation. *Biometrika*, 87(2):301–314, 2000.

[Bes86] Julian Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, 48(3):259–302, 1986.

[BH93] Daniel Barry and John A Hartigan. A Bayesian analysis for change point problems. *Journal of the American Statistical Association*, 88(421):309–319, 1993.

[BK04] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(9):1124–1137, 2004.

[BKL+09] Leif Boysen, Angela Kempe, Volkmar Liebscher, Axel Munk, and Olaf Wittich. Consistencies and rates of convergence of jump-penalized least squares estimators. *The Annals of Statistics*, 37(1):157–183, 2009.

[BM01] Lucien Birgé and Pascal Massart. Gaussian model selection. *Journal of the European Mathematical Society*, 3(3):203–268, 2001.

[BM07] Lucien Birgé and Pascal Massart. Minimal penalties for Gaussian model selection. *Probability Theory and Related Fields*, 138(1-2):33–73, 2007.

[BRT09] Peter J Bickel, Ya'acov Ritov, and Alexandre B Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.

[BVZ01] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on pattern analysis and machine intelligence*, 23(11):1222–1239, 2001.

[CDS01] Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.

[CL97] Antonin Chambolle and Pierre-Louis Lions. Image recovery via total variation minimization and related problems. *Numerische Mathematik*, 76(2):167–188, 1997.

[CZ64] Herman Chernoff and Shelemyahu Zacks. Estimating the current mean of a normal distribution which is subjected to changes in time. *The Annals of Mathematical Statistics*, 35(3):999–1018, 1964.

[DHL17] Arnak S Dalalyan, Mohamed Hebiri, and Johannes Lederer. On the prediction performance of the Lasso. *Bernoulli*, 23(1):552–581, 2017.

[DJ94] David L Donoho and Iain M Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455, 1994.

[DK01] P Laurie Davies and Arne Kovac. Local extremes, runs, strings and multiresolution. *The Annals of Statistics*, 29(1):1–48, 2001.

[Efr04] Bradley Efron. Large-scale simultaneous hypothesis testing: the choice of a null hypothesis. *Journal of the American Statistical Association*, 99(465):96–104, 2004.

[GBS08] Arpita Ghosh, Stephen Boyd, and Amin Saberi. Minimizing effective resistance of a graph. *SIAM review*, 50(1):37–66, 2008.

[GG84] Stuart Geman and Donald Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, 6:721–741, 1984.

[GO09] Tom Goldstein and Stanley Osher. The split Bregman method for L1-regularized problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343, 2009.

[GPS89] Dorothy M Greig, Bruce T Porteous, and Allan H Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, 51(2):271–279, 1989.

[Har16] Xiaoying Tian Harris. Prediction error after model search. *arXiv preprint arXiv:1610.06107*, 2016.

[HL13] Mohamed Hebiri and Johannes Lederer. How correlations influence lasso prediction. *IEEE Transactions on Information Theory*, 59(3):1846–1854, 2013.

[HLL10] Zaıd Harchaoui and Céline Lévy-Leduc. Multiple change-point estimation with a total variation penalty. *Journal of the American Statistical Association*, 105(492):1480–1493, 2010.

[HMPSST16] Oscar Hernan Madrid Padilla, James G Scott, James Sharpnack, and Ryan J Tibshirani. The DFS fused lasso: nearly optimal linear-time denoising over graphs and trees. *arXiv preprint arXiv:1608.03384*, 2016.

[Hoe10] Holger Hoefling. A path algorithm for the fused lasso signal approximator. *Journal of Computational and Graphical Statistics*, 19(4):984–1006, 2010.

[HR16] Jan-Christian Hütter and Philippe Rigollet. Optimal rates for total variation denoising. In *29th Annual Conference on Learning Theory*, pages 1115–1146, 2016.

[Joh15] Iain Johnstone. *Gaussian estimation: Sequence and wavelet models*. Available at statweb.stanford.edu/ imj/GE09-08-15.pdf, 2015.

[JSB+05] Brad Jackson, Jeffrey D Scargle, David Barnes, Sundararajan Arabhi, Alina Alt, Peter Gioumousis, Elyus Gwin, Paungkaew Sangtrakulcharoen, Linda Tan, and Tun Tao Tsai. An algorithm for optimal partitioning of data on an interval. *IEEE Signal Processing Letters*, 12(2):105–108, 2005.

[KFE12] Rebecca Killick, Paul Fearnhead, and Idris A Eckley. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500):1590–1598, 2012.

[KMP10] Ioannis Koutis, Gary L Miller, and Richard Peng. Approaching optimality for solving SDD linear systems. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 235–244. IEEE, 2010.

[KS11] Arne Kovac and Andrew DAC Smith. Nonparametric regression on a graph. *Journal of Computational and Graphical Statistics*, 20(2):432–447, 2011.

[KZ04] Vladimir Kolmogorov and Ramin Zabin. What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):147–159, 2004.

[L93] Lovász L. Random walks on graphs: A survey. *Combinatorics, Paul Erdös is eighty*, 2:1–46, 1993.

[LB12] Oren E Livne and Achi Brandt. Lean algebraic multigrid (LAMG): Fast graph Laplacian linear solver. *SIAM Journal on Scientific Computing*, 34(4):B499–B522, 2012.

[Leb05] Émilie Lebarbier. Detecting multiple change-points in the mean of Gaussian process by model selection. *Signal processing*, 85(4):717–736, 2005.

[LF97] Stephanie R Land and Jerome H Friedman. Variable fusion: A new adaptive signal regression method. Technical report, Technical Report 656, Department of Statistics, Carnegie Mellon University Pittsburgh, 1997.

[LSRT16] Kevin Lin, James Sharpnack, Alessandro Rinaldo, and Ryan J Tibshirani. Approximate recovery in changepoint problems, from ℓ_2 estimation error rates. *arXiv preprint arXiv:1606.06746*, 2016.

[MN00] Cristopher Moore and Mark EJ Newman. Epidemics and percolation in small-world networks. *Physical Review E*, 61(5):5678, 2000.

[MvdG97] Enno Mammen and Sara van de Geer. Locally adaptive regression splines. *The Annals of Statistics*, 25(1):387–413, 1997.

[OVLW04] Adam B Olshen, ES Venkatraman, Robert Lucito, and Michael Wigler. Circular binary segmentation for the analysis of array-based DNA copy number data. *Biostatistics*, 5(4):557–572, 2004.

[PXP00] Dzung L Pham, Chenyang Xu, and Jerry L Prince. Current methods in medical image segmentation 1. *Annual review of biomedical engineering*, 2(1):315–337, 2000.

[Rin09] Alessandro Rinaldo. Properties and refinements of the fused lasso. *The Annals of Statistics*, 37(5B):2922–2952, 2009.

[ROF92] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, 1992.

[RWY11] Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q-balls. *IEEE Transactions on Information Theory*, 57(10):6976–6994, 2011.

[Sha13] James Sharpnack. *Graph structured normal means inference*. PhD thesis, Carnegie Mellon University, 2013.

[SKS13] James L Sharpnack, Akshay Krishnamurthy, and Aarti Singh. Near-optimal anomaly detection in graphs using lovasz extended scan statistic. In *Advances in Neural Information Processing Systems*, pages 1959–1967, 2013.

[SRS12] James Sharpnack, Alessandro Rinaldo, and Aarti Singh. Sparsistency of the edge lasso over graphs. In *International Conference on Artificial Intelligence and Statistics*, pages 1028–1036, 2012.

[SS11] Daniel A Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.

[SSR13] James Sharpnack, Aarti Singh, and Alessandro Rinaldo. Detecting activations over graphs using spanning tree wavelet bases. In *International Conference on Artificial Intelligence and Statistics*, volume 16, pages 545–553, 2013.

[ST04] Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90. ACM, 2004.

[SWT16] Veeranjaneyulu Sadhanala, Yu-Xiang Wang, and Ryan J Tibshirani. Total variation classes beyond 1d: Minimax rates, and the limitations of linear smoothers. In *Advances in Neural Information Processing Systems*, pages 3513–3521, 2016.

[Tib96] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.

[TS15] Wesley Tansey and James G Scott. A fast and flexible algorithm for the graph-fused lasso. *arXiv preprint arXiv:1505.06475*, 2015.

[TSR+05] Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.

[TT11] Ryan J Tibshirani and Jonathan Taylor. The solution path of the generalized lasso. *The Annals of Statistics*, 39(3):1335–1371, 2011.

[TT15] Xiaoying Tian and Jonathan E Taylor. Selective inference with a randomized response. *arXiv preprint arXiv:1507.06739*, 2015.

[VDGB09] Sara A Van De Geer and Peter Bühlmann. On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.

[vdGL13] Sara van de Geer and Johannes Lederer. The lasso, correlated design, and improved oracle inequalities. In *From Probability to Statistics and Back: High-Dimensional Models and Processes–A Festschrift in Honor of Jon A. Wellner*, pages 303–316. Institute of Mathematical Statistics, 2013.

[WL02] Gerhard Winkler and Volkmar Liebscher. Smoothers for discontinuous signals. *Journal of Nonparametric Statistics*, 14(1-2):203–222, 2002.

[WSST16] Yu-Xiang Wang, James Sharpnack, Alex Smola, and Ryan J Tibshirani. Trend filtering on graphs. *Journal of Machine Learning Research*, 17(105):1–41, 2016.

[YA89] Yi-Ching Yao and Siu-Tong Au. Least-squares estimation of a step function. *Sankhyā: The Indian Journal of Statistics, Series A*, 51(3):370–381, 1989.

[Yao84] Yi-Ching Yao. Estimation of a noisy discrete-time step function: Bayes and empirical Bayes approaches. *The Annals of Statistics*, 12(3):1434–1447, 1984.

[Yao88] Yi-Ching Yao. Estimating the number of change-points via Schwarz' criterion. *Statistics & Probability Letters*, 6(3):181–189, 1988.

[ZBS01] Yongyue Zhang, Michael Brady, and Stephen Smith. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*, 20(1):45–57, 2001.

[ZWJ14] Yuchen Zhang, Martin J Wainwright, and Michael I Jordan. Lower bounds on the performance of polynomial-time algorithms for sparse linear regression. In *Conference on Learning Theory*, volume 35, pages 1–28, 2014.

[ZWJ15] Yuchen Zhang, Martin J Wainwright, and Michael I Jordan. Optimal prediction for sparse linear models? Lower bounds for coordinate-separable M-estimators. *arXiv preprint arXiv:1503.03188*, 2015.